

Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of CYBERATTACK CAPABILITIES

William A. Owens, Kenneth W. Dam, and Herbert S. Lin, *Editors*

Committee on Offensive Information Warfare

Computer Science and Telecommunications Board

Division on Engineering and Physical Sciences

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, N.W. Washington, DC 20001

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

Support for this project was provided by the MacArthur Foundation under award number 04-80965-000-GSS, the Microsoft Corporation under an unnumbered award, and the NRC Presidents' Committee under an unnumbered award.

Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the organizations that provided support for the project.

International Standard Book Number-13: 978-0-309-13850-5
International Standard Book Number-10: 0-309-13850-7
Library of Congress Control Number: 2009930416

Additional copies of this report are available from:

The National Academies Press
500 Fifth Street, N.W., Lockbox 285
Washington, DC 20055
(800) 624-6242
(202) 334-3313 (in the Washington metropolitan area)
Internet: <http://www.nap.edu>

Copyright 2009 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

THE NATIONAL ACADEMIES

Advisers to the Nation on Science, Engineering, and Medicine

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Ralph J. Cicerone is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Charles M. Vest is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Harvey V. Fineberg is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Ralph J. Cicerone and Dr. Charles M. Vest are chair and vice chair, respectively, of the National Research Council.

www.national-academies.org

COMMITTEE ON OFFENSIVE INFORMATION WARFARE

WILLIAM A. OWENS, AEA Holdings, Inc., *Co-chair*
KENNETH W. DAM, University of Chicago, *Co-chair*
THOMAS A. BERSON, Anagram Laboratories
GERHARD CASPER, Stanford University
DAVID D. CLARK, Massachusetts Institute of Technology
RICHARD L. GARWIN, IBM Fellow Emeritus
JACK L. GOLDSMITH III, Harvard Law School
CARL G. O'BERRY, The Boeing Company
JEROME H. SALTZER, Massachusetts Institute of Technology (retired)
MARK SEIDEN, MSB Associates
SARAH SEWALL, Harvard University
WALTER B. SLOCOMBE, Caplin & Drysdale
WILLIAM O. STUDEMANN, U.S. Navy (retired)
MICHAEL A. VATIS, Steptoe & Johnson LLP

Staff

HERBERT S. LIN, Study Director
KRISTEN BATCH, Associate Staff Officer (through August 2008)
TED SCHMITT, Consultant
JANICE SABUDA, Senior Project Assistant (through March 2008)
ERIC WHITAKER, Senior Project Assistant

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD

JOSEPH F. TRAUB, Columbia University, *Chair*
PRITHVIRAJ BANERJEE, Hewlett Packard Company
FREDERICK R. CHANG, University of Texas, Austin
WILLIAM DALLY, Stanford University
MARK E. DEAN, IBM Almaden Research Center
DEBORAH L. ESTRIN, University of California, Los Angeles
KEVIN C. KAHN, Intel Corporation
JAMES KAJIYA, Microsoft Corporation
RANDY H. KATZ, University of California, Berkeley
JOHN E. KELLY III, IBM Research
SARA KIESLER, Carnegie Mellon University
JON KLEINBERG, Cornell University
PETER LEE, Carnegie Mellon University
TERESA H. MENG, Stanford University
WILLIAM H. PRESS, University of Texas, Austin
PRABHAKAR RAGHAVAN, Yahoo! Research
DAVID E. SHAW, D.E. Shaw Research
ALFRED Z. SPECTOR, Google, Inc.
ROBERT F. SPROULL, Sun Microsystems, Inc.
PETER SZOLOVITS, Massachusetts Institute of Technology
ANDREW J. VITERBI, Viterbi Group, LLC
PETER WEINBERGER, Google, Inc.

JON EISENBERG, Director
RENEE HAWKINS, Financial and Administrative Manager
HERBERT S. LIN, Chief Scientist, CSTB
LYNETTE I. MILLETT, Senior Program Officer
NANCY GILLIS, Program Officer
ENITA A. WILLIAMS, Associate Program Officer
MORGAN R. MOTTO, Program Associate
SHENAE BRADLEY, Senior Program Assistant
ERIC WHITAKER, Senior Program Assistant

For more information on CSTB, see its website at <http://www.cstb.org>, write to CSTB, National Research Council, 500 Fifth Street, N.W., Washington, DC 20001, call (202) 334-2605, or e-mail CSTB at cstb@nas.edu.

Preface

Given the reality of a densely interconnected information society, much has been written about the possibility that adversaries of the United States such as terrorists or hostile nations might conduct very damaging cyberattacks against critical sectors of the U.S. economy and critical national infrastructure that depend on reliably functioning, secure computer systems and networks. For some years, the topic of cybersecurity has been an important part of the report portfolio of the National Research Council,¹ and a great deal of national attention has been given, in public, to the problem of how to protect U.S. information technology systems and networks against such attacks—that is, how to defend these systems and networks in both military and non-military contexts.² But, perhaps reflecting the common wisdom of the time, these efforts have focused almost exclusively on the cyberdefense side of the equation.

The possibility that the United States might choose to engage in cyberattacks to serve its own national interests—in cyberdefense as well

¹ An old but still quite relevant report on this topic is CSTB/National Research Council, *Computers at Risk*, National Academy Press, Washington, D.C., 1991; other relevant NRC reports include CSTB/NRC, *Trust in Cyberspace*, National Academy Press, Washington, D.C., 1999, and NRC, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

² See, for example, National Research Council, *Information Technology for Counterterrorism*, The National Academies Press, Washington, D.C., 2003; NRC, *Cybersecurity Today and Tomorrow: Pay Now or Pay Later*, The National Academies Press, Washington, D.C., 2002; and CSTB/NRC, *Realizing the Potential of C4I: Fundamental Challenges*, National Academy Press, Washington, D.C., 1998.

as in other areas—is rarely discussed in public. One recent public hint of U.S. government interest in the topic can be found in the still-classified Comprehensive National Cybersecurity Initiative (CNCI), which was adopted as national policy in January 2008 as part of National Security Presidential Directive 54/Homeland Security Presidential Directive 23 (NSPD-54/HSPD-23). According to the director of national intelligence in February 2009, “The CNCI addresses current cybersecurity threats, anticipates future threats and technologies, and develops a framework for creating in partnership with the private sector an environment that no longer favors cyber intruders over defenders. The CNCI includes defensive, *offensive* [emphasis added], education, research and development, and counterintelligence elements.”³ Press reports indicated that the CNCI involves 12 components designed to protect computer networks and systems and to improve information technology processes and policies.⁴ These components included a program to reduce the number of connections from federal agencies to external computer networks to 100 or fewer. The other 11 programs address intrusion detection; intrusion prevention; research and development; situational awareness (involving the coordination of information from all agencies to help secure cyber networks and systems); cyber counterintelligence; classified network security; cyber education and training; implementation of information security technologies; *deterrence strategies* [emphasis added]; global supply chain security; and public/private collaboration.

There is some public writing on the subject of cyberattack. Starting in the mid-1990s, the first papers on the topic emerged, many of them focusing on the legal issues involved in military uses of cyberattack.⁵ One of

³ Dennis Blair, Director of National Intelligence, Annual Threat Assessment of the Intelligence Community for the Senate Select Committee on Intelligence, February 12, 2009, available at <http://intelligence.senate.gov/090212/blair.pdf>.

⁴ See Jill R. Aitoro, “National Cyber Security Initiative Will Have a Dozen Parts,” *Government Executive*, August 1, 2008, available at http://www.nextgov.com/nextgov/ng_20080801_9053.php.

⁵ See, for example, Lawrence T. Greenberg, Seymour E. Goodman, and Kevin J. Soo Hoo, *Information Warfare and International Law*, National Defense University Press, 1998; Michael Schmitt, “Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework,” *Columbia Journal of Transnational Law* 37:885-937, 1999; Christopher C. Joyner and Catherine Lotrionte, “Information Warfare as International Coercion: Elements of a Legal Framework,” *European Journal of International Law* 12(5):825-865, 2001; Jason Barkham, “Information Warfare and International Law on the Use of Force,” *New York University International Law and Politics* 34:57-113, 2001; Davis Brown, “A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict,” *Harvard International Law Journal* 47(1):179-221, Winter 2006; Duncan B. Hollis, “New Tools, New Rules: International Law and Information Operations,” pp. 59-72 in *Ideas as Weapons: Influence and Perception in Modern Warfare*, G. David and T. McKeldin, eds., Potomac Books Inc., 2009.

the first studies to address the strategic implications of cyberattack was published by the RAND Corporation in 1996 (*Strategic Information Warfare: A New Face of War*).⁶ A later study covering the same topic in much more detail was published as *Strategic Warfare in Cyberspace*.⁷ A flurry of writing began to appear in the professional military literature in the late 1990s and early 2000s, but little or nothing can be found in this body of literature since around 2002 or 2003.

THIS STUDY—FOCUS, APPROACH, AND PURPOSE

Most of the writing to date has not brought together information technology experts who are knowledgeable in detail about what can and cannot be done from a technical standpoint with senior individuals who have policy experience, nor has it addressed the topic in an interdisciplinary manner that integrates expertise from the disciplines and fields that are relevant to the subject. The National Research Council undertook the present study (Box P.1) believing in the value of an integrated treatment that would help shed much-needed light on various important dimensions of cyberattack (and secondarily on the topic of cyberexploitation, a term that refers to the penetration of adversary computers and networks to obtain information for intelligence purposes). Such a treatment would provide a point of departure for others so that a broad variety of independent intellectual perspectives can be brought to bear on it.

The Committee on Offensive Information Warfare first met in July 2006 and five times subsequently. Its earlier meetings were devoted primarily to briefings on a variety of topics related to cyberattack, and later meetings were devoted primarily to committee deliberations.

The authoring committee did not receive classified information in the course of this study. What is sensitive about cyberattack is generally the fact of U.S. interest in a specific technology for cyberattack (rather than the nature of that technology itself); fragile and sensitive operational details that are not specific to the technologies themselves (e.g., the existence of a covert operative in a specific foreign country or a particular vulnerability); or capabilities and intentions of specific adversaries. None of these specific areas are particularly relevant to a study that focuses on the articulation of an intellectual framework for thinking about cyberattack.

It is important to delineate the scope of what this report does and

⁶ Roger C. Molander, Andrew S. Riddile, and Peter A. Wilson, *Strategic Information Warfare: A New Face of War*, National Defense Research Institute, RAND, Washington, D.C., 1996, available at http://www.rand.org/pubs/monograph_reports/2005/MR661.pdf.

⁷ Gregory J. Rattray, *Strategic Warfare in Cyberspace*, MIT Press, Cambridge, Mass., 2001.

BOX P.1 Statement of Task

The National Research Council will appoint an ad hoc committee to examine policy dimensions and legal/ethical implications of offensive information warfare, informed by expert perspectives on and knowledge of the underlying technologies. These policy dimensions include but are not limited to factors that differentiate between cyberattack as a law enforcement matter versus cyberattack as a national security matter, the extent to which the U.S. Department of Defense is constrained from acting in response to cyberattack of uncertain origin, appropriate definitions of concepts such as “force” or “armed attack” as they apply to different forms of offensive information warfare, the standards of proof required to establish the origin of a cyberattack, the nature and extent of actions that the United States may take unilaterally against a foreign cyberattacker, the possible utility of offensive information warfare as a mode of attack that is different from kinetic or physical attack, the nature and extent to which offensive information warfare may be a part of conventional military operations, and the extent to which a nation undertaking offensive information warfare may increase the likelihood that it would be attacked in response, either similarly or dissimilarly. Project products will be directed at policy makers and researchers, the former so that decision making can occur in a more informed manner and the latter so that other independent researchers will have a firm base on which to ground their own work.

does not do. This report does not provide a detailed explication of U.S. policy and practice regarding cyberattack or cyberexploitation, nor does it describe cyberattack/cyberexploitation capabilities available to the U.S. government. Instead, it provides a framework for understanding cyberattack that describes the basic technologies of cyberattack and cyberexploitation, articulates basic principles of what cyberattack and cyberexploitation might do, and discusses some of the policy goals that these actions might serve. It addresses some of the legal and ethical considerations that such uses might entail, and it suggests some analytical tools that might be useful for thinking about cyberattack from a policy perspective. It includes a number of findings and recommendations.

Just as other areas of national security have benefited from a vigorous public airing of issues, the authoring committee hopes that this report will stimulate debate and discussion on cyberattack as an instrument of national policy at the nexus of technology, policy, law, ethics, and national security both inside and outside government and thus bring to bear on these knotty issues the best intellectual thought and consideration.

A historical analogy might be drawn to the study of nuclear issues. In many ways, today’s state of affairs regarding public discourse on

cyberattack is analogous to the nuclear debate of 50 years ago. At that time, nuclear policy issues were veiled in secrecy, and there was little public debate about them. Herman Kahn's books (*On Thermonuclear War*, *Thinking the Unthinkable*) were the first that addressed in the open literature what it might mean to fight a nuclear war. These seminal pieces did much to raise the public profile of these issues and stimulated an enormous amount of subsequent work outside government that has had a real impact on nuclear policy.

From our perspective as the co-chairs of this study, the topic of cyberattack is so important across a multitude of national interests—not just defense or even just national security—that it deserves robust and open discussion and debate, both among thoughtful professionals in the policy, military, intelligence, law enforcement, and legal fields and among security practitioners in the private sector. But for such discussion and debate to be productive, they must be based on some common foundation of information about the topic at hand. Thus, the report's role in providing education and background is in our view its most important function.

It is because of the potential relevance of cyberattack across a broad spectrum of national interests that it was necessary to constitute a study committee whose members had very diverse backgrounds and many different perspectives on various aspects of cyberattack. The committee assembled for this project included individuals with expertise in networking, computer security, large-scale computer systems and architecture, national security and defense policy, intelligence and military operations, international law governing war and conflict, human rights, international relations and diplomacy, constitutional law and civil liberties, and domestic law enforcement as it relates to cybersecurity. Nonetheless, no one person had all of the necessary expertise across all relevant areas, and the committee expended considerable effort to bring all of its members to a common (if basic) level of understanding in these areas. The committee was by design highly heterogeneous and interdisciplinary, a characteristic intended to promote discussion and synergy among its members. As for the two co-chairs, one of us (Owens) has extensive military experience and the other (Dam) has extensive experience in foreign affairs—and both of us have served in the private sector.

We hope that a second function of this report is to help establish the awareness needed in the elected government (executive and legislative) for making good decisions and providing proper oversight of capabilities for cyberattack. As the report points out, the U.S. government does not appear to be well organized to manage these capabilities, either in an employment sense (when and under what circumstances a particular kind of cyberattack should be used) or in an acquisition sense (how to obtain capabilities for cyberattack). Many of the report's findings and recom-

mendations focus on taking some first steps for better organization of the government in this regard.

How will the presentation and analysis in this report endure over time? The question is natural given the inevitability of changes in the technological and global environment. Based on historical experience, it is highly likely that in a decade the technological substrate underlying information technology applications will be very different by one, two, or three orders of magnitude in a number of dimensions—processor power, cost, bandwidth, storage, and so on. Deployments of information technology will be more pervasive. Connectivity among individuals, embedded computers, and organizations is likely to increase dramatically, and myriad applications that are entirely unimagined now will be commonplace. The importance of the Internet (or its follow-on) to society will be greater as well.

The global environment is also likely to be substantially different, although *how* it will be different cannot be predicted in the same way that Moore's law predicts circuit densities. Many analysts of international affairs predict a rise in the significance of actors not tied or only loosely tied to nation-states, or of adversaries that do not share U.S. or Western values or legal traditions with respect to the conduct of conflict.

Few portions of the report are tied explicitly to current technologies, although a genuine breakthrough in technologies to support non-cooperative high-precision attribution of attacks to bad actors in cyberspace would have significant implications for several findings in this report. A rise in the non-state actor cyberthreat would be significant as well—and although the committee has attempted to grapple with that issue in its analysis, it would be the first to admit that a great deal more work is needed in that area. Thus, it is the committee's hope that the framework established in this report for understanding cyberattack will endure for some time to come.

As this report goes into final publication in July 2009, a number of significant events have occurred. For example, on May 29, 2009, the Obama White House released its 60-day review of cybersecurity policy (*Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure*⁸), a document that is essentially silent on the offensive dimension of cybersecurity. On June 23, Secretary of Defense Robert Gates directed the establishment of the U.S. Cyber Command, a sub-unified command subordinate to U.S. Strategic Command and responsible for military cyberspace operations.⁹

⁸See http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf.

⁹Siobhan Gorman and Yochi Dreazen, "Military Command Is Created for Cyber Security," *Wall Street Journal*, June 24, 2009, available at <http://online.wsj.com/article/SB124579956278644449.html>.

ACKNOWLEDGMENTS

This study could not have been undertaken without the financial support of the MacArthur Foundation and the Microsoft Corporation, both of which recognized the potential legal, policy, and ethical significance of new technologies for cyberattack. The National Research Council itself also provided some funding for this project.

The complexity of the issues explored in this report meant that the committee had much to learn from its briefers. The committee is grateful to many individuals:

- For briefings on cyberattack technologies, Steven Bellovin of Columbia University and William Howard, independent consultant;
- For briefings on operational dimensions of cyberattack, Patrick D. Allen of General Dynamics Advanced Information Systems, Lt. Gen. Bill Donahue, U.S. Air Force (retired), and Sam Gardiner, U.S. Air Force (retired);
- For briefings on the various legal dimensions of cyberattack and cyberexploitation, Thomas Wingfield of the Potomac Institute, LTC Eric Jensen of the Office of the Judge Advocate General, U.S. Army, Joe Dhillon of the McGeorge School of Law at the University of the Pacific, Jeff Smith (former CIA General Counsel), Jim Dempsey of the Center for Democracy and Technology, Richard Salgado of Yahoo!, Eugene Volokh of the UCLA School of Law, and Robert Weisberg and Helen Stacy of the Stanford University Law School;
- For briefings on the ethics of cyberattack, Jeff McMahan of Rutgers University and Father J. Bryan Hehir of Harvard University;
- For briefings on current DOD perspectives on cyberattack, Admiral Elizabeth Hight of the JTFGNO, LTC Forrest Hare of the U.S. Air Force, and Dr. Linton Wells of the Department of Defense;
- For briefings on policy issues regarding non-lethal weapons, David Koplou of Georgetown University;
- For briefings on private sector perspectives, Rod Wallace of Nortel, Milo Medin of M2Z Networks, Jeffrey I. Schiller of MIT, and Naveen Jain of Intelius;
- For a briefing on deterrence theory as it might be applied to cyber-attack, Thomas Schelling of the University of Maryland; and
- For a variety of independent perspectives on the subject of cyber-attack, Stephen Cambone (former Undersecretary of Defense for Intelligence), James N. Miller, Jr. (former Deputy Assistant Secretary of Defense for Requirements, Plans, and Counterproliferation), Dan Kuehl of the National Defense University, Stuart Starr of the Center for Technology and National Security Policy at the National Defense University, K.A. Taipale of the Center for Advanced Studies in Science and Technology

Policy, Neal Pollard of Georgetown University and the National Counterterrorism Center, and Dorothy E. Denning of the Naval Postgraduate School.

Throughout the study, the committee's complex and challenging technical and political discussions encompassed a wide range of thought and perspectives offered by those who appeared before the committee, the committee itself, and participants in the review process. In addition, from the CSTB staff, we thank Ted Schmitt, Kristen Batch, and David Padgham for substantial research assistance, and Janice Sabuda and Eric Whitaker for administrative support. Lastly, this report would not have been possible without the leadership and stamina of Herb Lin, whose organizational skills, leadership of the staff, and thoughtful, complete agendas for committee discussion exceeded the excellent standards established by the National Academies. Both of us are indebted to Herb for his counsel, his policy and technical knowledge, his ability to find "just the right wording" for contentious areas, and for the character and chemistry he has shown us in our personal dealings. The National Research Council is fortunate to have leaders of Herb Lin's quality. This study and we co-chairs have benefited greatly from his deep involvement.

Kenneth W. Dam and William A. Owens, *Co-chairs*
Committee on Offensive Information Warfare

Acknowledgment of Reviewers

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the National Research Council's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

Matt Blaze, University of Pennsylvania,
W. Earl Boebert, Sandia National Laboratories (retired),
Lewis M. Branscomb, Independent Consultant, La Jolla, California,
Jogindar (Joe) Dhillon, State of California,
Stephen Dycus, Vermont Law School,
Michael Froomkin, University of Miami School of Law,
Dan Geer, Geer Risk Services,
Ronald Lee, Arnold and Porter,
Martin Libicki, RAND Corporation,
James McCarthy, USAF Academy,
John McLaughlin, Johns Hopkins University,
Richard Mies, SAIC,
Gregory Rattray, Independent Consultant, San Antonio, Texas,

Abe Sofaer, Stanford University,
Eugene Spafford, Purdue University,
Phil Venables, Goldman Sachs & Co.,
Peter Weinberger, Google, Inc., and
Marc J. Zwillinger, Sonnenschein Nath & Rosenthal.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations, nor did they see the final draft of the report before its release. The review of this report was overseen by William H. Press, University of Texas at Austin, and Eugene Volokh, University of California at Los Angeles. Appointed by the National Research Council, they were responsible for making certain that an independent examination of this report was carried out in accordance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

Contents

SYNOPSIS	1
1 OVERVIEW, FINDINGS, AND RECOMMENDATIONS	9
1.1 What Is Cyberattack and Why Is It Important?, 9	
1.2 Focus of and Motivation for This Report, 12	
1.3 Cyberattack in the Context of an Information Strategy for the United States, 17	
1.4 Important Characteristics of Cyberattack and Cyberexploitation, 19	
1.5 Illustrative Applications of Cyberattack, 21	
1.6 The Legal Framework Governing Cyberattack, 21	
1.7 The Dynamics of Cyberconflict, 22	
1.8 Findings, 24	
1.8.1 Technologies as Instruments of U.S. National Policy, 24	
1.8.2 Overarching Findings, 25	
1.8.3 Legal and Ethical Findings, 31	
1.8.4 Policy Findings, 39	
1.8.5 Technical and Operational Findings, 43	
1.8.6 Organizational Findings, 53	
1.9 Recommendations, 56	
1.9.1 Fostering a National Debate on Cyberattack, 57	
1.9.2 Organizing the Decision-Making Apparatus of the U.S. Government for Cyberattack, 62	
1.9.3 Supporting Cyberattack Capabilities and Policy, 66	

- 1.9.4 Developing New Knowledge and Insight into a New Domain of Conflict, 73
- 1.10 Conclusion, 75

**PART I
FRAMING AND BASIC TECHNOLOGY**

- 2 TECHNICAL AND OPERATIONAL CONSIDERATIONS IN CYBERATTACK AND CYBEREXPLOITATION 79
 - 2.1 Important Characteristics of Cyberattack and Cyberexploitation, 80
 - 2.2 The Basic Technology of Cyberattack, 82
 - 2.2.1 Information Technology and Infrastructure, 82
 - 2.2.2 Vulnerability, Access, and Payload, 83
 - 2.2.3 Scale and Precision, 89
 - 2.2.4 Critical Periods of Cyberattack, 89
 - 2.2.5 Approaches for Cyberattack, 91
 - 2.2.6 Propagating a Large-Scale Cyber Offensive Action, 106
 - 2.2.7 Economics, 108
 - 2.3 Operational Considerations, 110
 - 2.3.1 The Effects of Cyberattack, 110
 - 2.3.2 Possible Objectives of Cyberattack, 114
 - 2.3.3 Target Identification, 116
 - 2.3.4 Intelligence Requirements and Preparation, 118
 - 2.3.5 Effects Prediction and Damage Assessment, 121
 - 2.3.6 Complexity, Information Requirements, and Uncertainty, 126
 - 2.3.7 Rules of Engagement, 128
 - 2.3.8 Command and Control, 129
 - 2.3.9 Coordination of Cyberattack Activities with Other Institutional Entities, 132
 - 2.3.10 A Rapidly Changing and Changeable Technology and Operational Environment for Cyberattack, 133
 - 2.4 Characterizing an Incoming Cyberattack, 134
 - 2.4.1 Tactical Warning and Attack Assessment, 135
 - 2.4.2 Attribution, 138
 - 2.4.3 Intent, 141
 - 2.5 Active Defense for Neutralization as a Partially Worked Example, 142

- 2.6 Technical and Operational Considerations for Cyberexploitation, 149
 - 2.6.1 Technical Similarities in and Differences Between Cyberattack and Cyberexploitation, 149
 - 2.6.2 Possible Objectives of Cyberexploitation, 150
 - 2.6.3 Approaches for Cyberexploitation, 152
 - 2.6.4 Some Operational Considerations for Cyberexploitation, 153
- 2.7 Historical Precedents and Lessons, 156

PART II

MISSION AND INSTITUTIONAL PERSPECTIVES

- 3 A MILITARY PERSPECTIVE ON CYBERATTACK 161
 - 3.1 U.S. Military Doctrine and Cyberattack, 161
 - 3.2 Department of Defense Organization for Cyberattack, 165
 - 3.3 Rules of Engagement, 167
 - 3.4 Some Historical Perspective, 171
 - 3.5 Cyberattack in Support of Military Operations—Some Hypothetical Examples, 177
 - 3.5.1 Cyberattack in Support of Defense, Exploitation, and Other Information Operations, 177
 - 3.5.2 Cyberattack in Support of Traditional Military Operations, 179
 - 3.5.3 Cyberattack in Support of Other Operations, 180
 - 3.6 Operational Planning, 182
 - 3.7 Human Capital and Resources, 184
 - 3.8 Weapons Systems Acquisition, 186
- 4 AN INTELLIGENCE COMMUNITY PERSPECTIVE ON CYBERATTACK AND CYBEREXPLOITATION 188
 - 4.1 Intelligence Collection and Analysis, 188
 - 4.1.1 Governing Principles, 188
 - 4.1.2 How Cyberexploitation Might Be Used to Support Intelligence Collection, 190
 - 4.2 Covert Action, 193
 - 4.2.1 Governing Principles, 193
 - 4.2.2 How Cyberattack Might Be Used in Covert Action, 195
 - 4.3 Possible Intelligence Community Interest in Cyberattack and Cyberexploitation, 198

5	PERSPECTIVES ON CYBERATTACK OUTSIDE NATIONAL SECURITY	200
5.1	Cyberattack and Domestic Law Enforcement, 200	
5.2	Threat Neutralization in the Private Sector, 202	
5.2.1	Possible Response Options for Private Parties Targeted by Cyberattack, 202	
5.2.2	Self-defense by Private Parties, 204	
5.2.3	Regulating Self-defense by Private Parties, 208	
5.2.4	Negative Ramifications of Self-defense by Private Parties, 210	
5.3	Cyberexploitation in the Private Sector, 212	
5.4	Threat Neutralization on Behalf of Non-military Government Agencies, 213	
6	DECISION MAKING AND OVERSIGHT	214
6.1	Executive Branch, 214	
6.1.1	Declaratory Policy, 215	
6.1.2	Acquisition Policy, 220	
6.1.3	Employment Policy, 223	
6.1.4	Operational Oversight, 229	
6.2	Legislative Branch, 232	
6.2.1	Warmaking Powers, 232	
6.2.2	Budget, 234	
6.2.3	Oversight (and Notification), 235	

PART III
INTELLECTUAL TOOLS FOR UNDERSTANDING AND
THINKING ABOUT CYBERATTACK

7	LEGAL AND ETHICAL PERSPECTIVES ON CYBERATTACK	239
7.1	The Basic Framework, 239	
7.2	International Law, 241	
7.2.1	The Law of Armed Conflict, 241	
7.2.2	Applying the Law of Armed Conflict to Cyberattack, 250	
7.2.3	International Law and Non-state Actors, 273	
7.2.4	The Convention on Cybercrime, 277	
7.2.5	Human Rights Law, 281	
7.2.6	Reciprocity, 282	

7.3	Domestic Law, 282	
7.3.1	Covert Action and Military Activity, 283	
7.3.2	Title III and the Foreign Intelligence Surveillance Act, 286	
7.3.3	Posse Comitatus, 288	
7.3.4	The Computer Fraud and Abuse Act and Other Federal Law, 288	
7.3.5	The War Powers Resolution, 290	
7.3.6	Executive Order 12333 (United States Intelligence Activities), 290	
7.4	Foreign Domestic Law, 292	
8	INSIGHTS FROM RELATED AREAS	293
8.1	Nuclear Weapons and Nuclear War, 293	
8.2	Space, 296	
8.3	Biological Weapons, 297	
8.4	Non-lethal Weapons, 299	
9	SPECULATIONS ON THE DYNAMICS OF CYBERCONFLICT	302
9.1	Deterrence and Cyberconflict, 302	
9.2	Escalatory Dynamics of Cyberconflict Between Nation-States, 306	
9.2.1	Crisis Stability, 306	
9.2.2	Escalation Control and Management, 308	
9.2.3	Complications Introduced by Patriotic Hackers, 310	
9.2.4	Incentives for Self-restraint in Escalation, 311	
9.2.5	Termination of Cyberconflict, 311	
9.2.6	The Role of Transparency, 312	
9.2.7	Catalytic Cyberconflict, 312	
9.3	Cyberconflict Between the United States and Non-state Actors, 313	
9.4	The Political Side of Escalation, 315	
10	ALTERNATIVE FUTURES	318
10.1	Regulatory Regimes—Basic Principles, 318	
10.2	Regulatory Regimes for Cyberattack, 321	
10.2.1	Direct Approaches Based on Traditional Arms Control, 321	
10.2.2	Indirect Approaches Based on Regulation of Non-military Domains, 326	
10.3	Foreign Perspectives on Cyberattack, 328	

APPENDIXES

A	Biographies of Committee Members and Staff	337
B	Meeting Participants and Other Contributors	348
C	Illustrative Criminal Cyberattacks	350
D	Views on the Use of Force in Cyberspace	356
E	Technical Vulnerabilities Targeted by Cyber Offensive Actions	360

Synopsis

This synopsis is intended to provide the reader with a sense of what the report contains. However, it is necessarily incomplete, and it omits any mention of many significant topics contained in the main body of the report.

UNDERSTANDING CYBERATTACK

What Is Cyberattack?

Cyberattack refers to deliberate actions to alter, disrupt, deceive, degrade, or destroy computer systems or networks or the information and/or programs resident in or transiting these systems or networks. The U.S. armed forces are actively preparing to engage in cyberattacks, perhaps in concert with other information warfare means and/or with kinetic attacks, and may have done so in the past. Domestic law enforcement agencies also engage in cyberattack when they jam cell phone networks in order to prevent the detonation of improvised explosive devices.

Such matters pose some very important issues that relate to technology, policy, law, and ethics. This report provides an intellectual framework for thinking about cyberattack and understanding these issues.

A first point is that cyberattack must be clearly distinguished from cyberexploitation, which is an intelligence-gathering activity rather than a destructive activity. Although much of the technology underlying cyberexploitation is similar to that of cyberattack, cyberattack and cyber-

exploitation are conducted for entirely different purposes. (This contrast is relevant to much of the public debate using the term “cyberattack,” which in common usage often lumps both attack and exploitation under the “attack” label.)

Second, weapons for cyberattack have a number of characteristics that differentiate them from traditional kinetic weapons. Compared to kinetic weapons, many weapons for cyberattack:

- Are easy to use with high degrees of anonymity and with plausible deniability, making them well suited for covert operations and for instigating conflict between other parties;
- Are more uncertain in the outcomes they produce, making it difficult to estimate deliberate and collateral damage; and
- Involve a much larger range of options and possible outcomes, and may operate on time scales ranging from tenths of a second to years, and at spatial scales anywhere from “concentrated in a facility next door” to globally dispersed.

Third, cyberattack as a mode of conflict raises many operational issues. For example, given that any large nation experiences cyberattacks continuously, how will the United States know it is the subject of a cyberattack deliberately launched by an adversary government? There is also a further tension between a policy need for rapid response and the technical reality that attribution is a time-consuming task. Shortening the time for investigation may well increase the likelihood of errors being made in a response (e.g., responding against the wrong machine or launching a response that has large unintended effects).

Illustrative Applications of Cyberattack

Cyberattack can support military operations. For example, a cyberattack could disrupt adversary command, control, and communications; suppress air defenses; degrade smart munitions and platforms; or attack warfighting or warmaking infrastructure (the defense industrial base). Cyberattack might be used to augment or to enable some other kinetic attack to succeed, or to defend a friendly computer system or network by neutralizing the source of a cyberattack conducted against it.

Cyberattack can also support covert action, which is designed to influence governments, events, organizations, or persons in support of foreign policy in a manner that is not necessarily attributable to the U.S. government. The range of possible cyberattack options is very large, and so cyberattack-based covert action might be used, for example, to

influence an election, instigate conflict between political factions, harass disfavored leaders or entities, or divert money.

Illustrative Applications of Cyberexploitation

For intelligence gathering, cyberexploitation of an adversary's computer systems might yield valuable information. For example, U.S. intelligence agencies might learn useful information about an adversary's intentions and capabilities from a penetration of its classified government networks. Alternatively, they might obtain useful economic information from penetrating the computer systems of a competing nation's major industrial firms.

The Legal Framework Governing Cyberattack

In the committee's view, the essential framework for the legal analysis of cyberattack is based on the principle that notions related to "use of force" and "armed attack" (terms of special relevance to the Charter of the United Nations) should be judged primarily by the effects of an action rather than its modality. That is, the fact that an attack is carried out through the use of cyberweapons rather than kinetic weapons is far less significant than the effects that result from such use, where "effects" are understood to include both direct and indirect effects.

Furthermore, the committee believes that the principles of the law of armed conflict (LOAC) and the Charter of the United Nations—including both law governing the legality of going to war (*jus ad bellum*) and law governing behavior during war (*jus in bello*)—do apply to cyberattack, although new analytical work may be needed to understand how these principles do or should apply to cyberweapons. That is, some types of cyberattack are difficult to analyze within the traditional LOAC structure. Among the more problematic cases are the following:

- The presumption of nation-to-nation conflict between national military forces,
- The exception for espionage, and
- The emphasis on notions of territorial integrity.

The Dynamics of Cyberconflict

The escalatory dynamics of armed conflict are thought to be understood as the result of many years of thinking about the subject, but the dynamics of cyberconflict are poorly understood. This report speculates on some of the factors that might influence the evolution of a cyberconflict.

For major nation-states with significant capabilities for kinetic attack and cyberattack at their disposal, among the important issues regarding the dynamics of cyberconflict are the following:

- Crisis stability (preventing a serious cyberconflict from breaking out),
- Preventing a cyberconflict from escalating to physical space, and
- Knowing when a cyberconflict has been terminated.

Matters can be further complicated by the presence of non-state actors, such as cyberterrorists, patriotic hackers, and criminal groups. Perhaps the most important complication relates to identification of the appropriate party against which action might be taken and the related availability of cyber and/or kinetic targets whose destruction might cause pain or meaningful damage to the terrorist or criminal group.

FINDINGS

Cyberattack is an important capability for the United States to maintain, but at the same time the acquisition and use of such capabilities raise many questions and issues, as described below.

Overarching Findings

1. The policy and organizational issues raised by U.S. acquisition and use of cyberattack are significant across a broad range of conflict scenarios, from small skirmishes with minor actors on the international stage to all-out conflicts with adversaries capable of employing weapons of mass destruction.

2. The availability of cyberattack technologies for national purposes greatly expands the range of options available to U.S. policy makers as well as to policy makers of other nations.

3. Today's policy and legal framework for guiding and regulating the U.S. use of cyberattack is ill-formed, undeveloped, and highly uncertain.

4. Secrecy has impeded widespread understanding and debate about the nature and implications of U.S. cyberattack.

5. The consequences of a cyberattack may be both direct and indirect, and in some cases of interest, the indirect consequences of a cyberattack can far outweigh the direct consequences.

Legal and Ethical Findings

6. The conceptual framework that underpins the UN Charter on the use of force and armed attack and today's law of armed conflict provides a reasonable starting point for an international legal regime to govern cyberattack. However, those legal constructs fail to account for non-state actors and for the technical characteristics of some cyberattacks.

7. In today's security environment, private parties have few useful alternatives for responding to a severe cyberattack that arrives over a network such as the Internet.

8. Cyberattack poses challenges to existing ethical and human rights regimes.

Policy Findings

9. Enduring unilateral dominance in cyberspace is neither realistic nor achievable by the United States.

10. The United States has much to lose from unrestrained cyberattack capabilities that are proliferated worldwide.

11. Deterrence of cyberattacks by the threat of in-kind response has limited applicability.

12. Options for responding to cyberattacks on the United States span a broad range and include a mix of dynamic changes in defensive postures, law enforcement actions, diplomacy, cyberattacks, and kinetic attacks.

Technical and Operational Findings

13. For many kinds of information technology infrastructure targets, the ease of cyberattack is increasing rather than decreasing.

14. Although the actual cyberattack capabilities of the United States are highly classified, they are at least as powerful as those demonstrated by the most sophisticated cyberattacks perpetrated by cybercriminals and are likely more powerful.

15. As is true for air, sea, land, and space operations, the defensive or offensive intent motivating cyber operations in any given instance may be difficult to infer.

16. Certain cyberattacks undertaken by the United States are likely to have significant operational implications for the U.S. private sector.

17. If and when the United States decides to launch a cyberattack, significant coordination among allied nations and a wide range of public and private entities may be necessary, depending on the scope and nature of the cyberattack in question.

18. The outcomes of many kinds of cyberattack are likely to be more uncertain than outcomes for other kinds of attack.

19. Early use of cyberattack may be easy to contemplate in a pre-conflict situation, and so a greater degree of operational oversight for cyberattack may be needed compared to that for the use of other options.

20. Developing appropriate rules of engagement for the use of cyberweapons is very difficult.

Organizational Findings

21. Both the decision-making apparatus for cyberattack and the oversight mechanisms for that apparatus are inadequate today.

22. The U.S. Congress has a substantial role to play in authorizing the use of military force, but the contours of that authority and the circumstances under which authorization is necessary are at least as uncertain for cyberattack as for the use of other weapons.

RECOMMENDATIONS

Fostering a National Debate on Cyberattack

1. The United States should establish a public national policy regarding cyberattack for all sectors of government, including but not necessarily limited to the Departments of Defense, State, Homeland Security, Treasury, and Commerce; the intelligence community; and law enforcement. The senior leadership of these organizations should be involved in formulating this national policy.

2. The U.S. government should conduct a broad, unclassified national debate and discussion about cyberattack policy, ensuring that all parties—particularly Congress, the professional military, and the intelligence agencies—are involved in discussions and are familiar with the issues.

3. The U.S. government should work to find common ground with other nations regarding cyberattack. Such common ground should include better mutual understanding regarding various national views of cyberattack, as well as measures to promote transparency and confidence building.

Organizing the Decision-Making Apparatus of the U.S. Government for Cyberattack

4. The U.S. government should have a clear, transparent, and inclusive decision-making structure in place to decide how, when, and why a cyberattack will be conducted.

5. The U.S. government should provide a periodic accounting of cyberattacks undertaken by the U.S. armed forces, federal law enforcement agencies, intelligence agencies, and any other agencies with authorities to conduct such attacks in sufficient detail to provide decision makers with a more comprehensive understanding of these activities. Such a periodic accounting should be made available both to senior decision makers in the executive branch and to the appropriate congressional leaders and committees.

Supporting Cyberattack Capabilities and Policy

6. U.S. policy makers should judge the policy, legal, and ethical significance of launching a cyberattack largely on the basis of both its likely direct effects and its indirect effects.

7. U.S. policy makers should apply the moral and ethical principles underlying the law of armed conflict to cyberattack even in situations that fall short of actual armed conflict.

8. The United States should maintain and acquire effective cyberattack capabilities. Advances in capabilities should be continually factored into policy development, and a comprehensive budget accounting for research, development, testing, and evaluation relevant to cyberattack should be available to appropriate decision makers in the executive and legislative branches.

9. The U.S. government should ensure that there are sufficient levels of personnel trained in all dimensions of cyberattack, and that the senior leaders of government have more than a nodding acquaintance with such issues.

10. The U.S. government should consider the establishment of a government-based institutional structure through which selected private sector entities can seek immediate relief if they are the victims of cyberattack.

Developing New Knowledge and Insight into a New Domain of Conflict

11. The U.S. government should conduct high-level wargaming exercises to understand the dynamics and potential consequences of cyberconflict.

12. Foundations and government research funders should support academic and think-tank inquiry into cyberconflict, just as they have supported similar work on issues related to nuclear, biological, and chemical weapons.

Overview, Findings, and Recommendations

1.1 WHAT IS CYBERATTACK AND WHY IS IT IMPORTANT?

It is now broadly accepted that nations are becoming ever more dependent on information and information technology. Companies and organizations rely on computers for diverse business processes ranging from payroll and accounting, to the tracking of inventory and sales, to support for research and development (R&D). Food, water, and energy distribution rely on computers and networks at every stage, as do transportation, health care, and financial services.

The same dependence also increasingly applies to the military. Modern military forces use weapons that are computer-controlled. Even more importantly, the movements and actions of military forces are increasingly coordinated through computer-based networks that allow information and common pictures of the battlefield to be shared. Logistics are entirely dependent on computer-based scheduling and optimization.

Even terrorists rely on information technology. Although the weapons of terrorists are generally low-tech, their use of the Internet and information technology for recruitment, training, and communications is often highly sophisticated.

Given the importance of information technology to many societal functions, it is not surprising that there has been much public debate about cybersecurity (i.e., protection of information technology systems and networks and the programs and information within them from hostile actions) and about how the United States might improve its cybersecurity posture in the face of hostile actions perpetrated by an adversary,

such as a terrorist group, criminals, or another nation. Although in many other domains, security has always had both defensive and attack components, cybersecurity has been somewhat anomalous, in the sense that its purely defensive side has been the primary focus of attention over the years. But, in fact, it is possible to imagine that cyberattacks might be used to support cyber defensive objectives. It is further possible to imagine that cyberattack would naturally be part of a robust U.S. military posture.

The possibility that the United States might choose to engage in cyberattacks to serve its own national interests is, however, rarely discussed in public. For the record, the U.S. government has acknowledged that it has an interest in such capabilities as a possible instrument of national policy,¹ but this is virtually all that it acknowledges publicly. At least one press report has indicated the existence of a still-classified National Security Presidential Directive, NSPD 16, issued in July 2002, that reportedly ordered the U.S. government to develop national-level guidance for determining when and how the United States would launch cyberattacks against enemy computer networks.² The *National Strategy to Secure Cyberspace*, published in February 2003, is entirely silent about an offensive component to U.S. cybersecurity efforts.³

In practice, hostile actions against a computer system or network can take two forms. One form is destructive in nature—the action is taken to harm the system or network and render it less functional or useful than before the action was taken. An example of such a hostile action is erasure by a computer virus of the hard disk of any computer that it infects. The second form is non-destructive—the action is taken to extract from a system or network information that would otherwise be kept confidential. Actions of this second form are usually clandestine, conducted with the smallest possible interventions that still allow extraction of the information sought. Such an action is exemplified by a computer virus that searches the hard disk of any infected computer and e-mails to the hostile party all files containing a credit card number.

Collectively, both forms of hostile action are termed “cyber offensive operations,” or simply, “cyber offense.” In this report, because the distinction between them is often important, the two forms of hostile action are given individual designators and somewhat expanded definitions:

- *Cyberattack* refers to the use of deliberate actions—perhaps over an extended period of time—to alter, disrupt, deceive, degrade, or destroy

¹ *An Assessment of International Legal Issues in Information Operations*, 2nd edition, Department of Defense, Office of General Counsel, November 1999.

² Bradley Graham, “Bush Orders Guidelines for Cyber-Warfare,” *Washington Post*, February 7, 2003, p. A01.

³ See http://www.dhs.gov/xlibrary/assets/National_Cyberspace_Strategy.pdf.

adversary computer systems or networks or the information and/or programs resident in or transiting these systems or networks.⁴ Such effects on adversary systems and networks may also have indirect effects on entities coupled to or reliant on them. A cyberattack seeks to cause adversary computer systems and networks to be unavailable or untrustworthy and therefore less useful to the adversary. Furthermore, because so many different kinds of cyberattack are possible, the term “cyberattack” should be understood as a statement about a methodology for action—and that alone—rather than as a statement about the scale of the action’s effect.

- *Cyberexploitation* refers to the use of cyber offensive actions—perhaps over an extended period of time—to support the goals and missions of the party conducting the exploitation, usually for the purpose of obtaining information resident on or transiting through an adversary’s computer systems or networks. Cyberexploitations do not seek to disturb the normal functioning of a computer system or network from the user’s point of view—indeed, the best cyberexploitation is one that such a user never notices.

Box 1.1 summarizes important distinctions between cyberattacks and cyberexploitations. The committee recognizes that analysts and commentators have used a variety of different terms that are closely related to what this report calls cyberattack (Box 1.2).

For purposes of this report, cyberattacks do not include kinetic actions taken against computers or networks using cruise missiles, sledgehammers, or satchel charges. But in practice, the destruction of or damage to an adversary computer or network could be accomplished by kinetic as well as cyber actions. Thus, as acknowledged by the Department of Defense,⁵ a planner contemplating the destruction of an adversary computer or network should think about both cyberattack and kinetic attack options. This report also does not consider the use of electromagnetic pulse (EMP) attacks. EMP attacks typically refer to non-selective attacks on electronics and electrical components on a large scale, although a tactical EMP weapon intended to selectively target such components on a small scale is possible to imagine.⁶

⁴ An adversary computer or network may not necessarily be owned and operated by the adversary—it may simply support or be used by the adversary.

⁵ “DoD will conduct kinetic missions to preserve freedom of action and strategic advantage in cyberspace. Kinetic actions can be either offensive or defensive and used in conjunction with other mission areas to achieve optimal military effects.” See Department of Defense, *National Military Strategy for Cyberspace Operations*, 2006, available at www.dod.mil/pubs/foi/ojcs/07-F-2105doc1.pdf.

⁶ For a comprehensive description of the threat from EMP attacks, see *Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack*, available at http://www.globalsecurity.org/wmd/library/congress/2004_r/04-07-22emp.pdf.

BOX 1.1 Cyberattack Versus Cyberexploitation

Terms ¹	Cyberattack, attack, computer network attack	Cyberexploitation, intelligence, exploitation, computer network exploitation
Approach and intent	Degrade, disrupt, deny, destroy attacked infrastructure and systems/networks	Conduct smallest intervention consistent with desired operations
Primary relevant domestic law	U.S. Code Title 10 authorities and restrictions ²	U.S. Code Title 50 authorities and restrictions
Operational agency	U.S. Strategic Command, Joint Functional Combatant Command for Network Warfare	National Security Agency
Main advocate in the U.S. government to date	U.S. Air Force	Director of National Intelligence
Interactions with tactical military operations	Based on explicit inclusion in battle plans	Based on intelligence reporting
Characterization of personnel	Warfighters	Intelligence community

¹ Discussion of these terms and concepts can be found in Chapters 2, 3, and 4.

² Covert action involving cyberattack would fall under Title 50 authorities.

1.2 FOCUS OF AND MOTIVATION FOR THIS REPORT

This report of the Committee on Offensive Information Warfare focuses primarily on the policy consequences and legal and ethical implications of U.S. acquisition and use of cyberattack, and secondarily (and only when necessary) on cyberexploitation. There are two reasons that a report on cyberattack necessarily touches on cyberexploitation. First, cyberattack and cyberexploitation are closely related from a technical point of view.

Second, because of such similarities a nation that is the target of a cyberexploitation might misinterpret it as being a cyberattack—a possibility that U.S. policy makers must take into account in deciding whether to conduct a cyberexploitation. Nevertheless, the policy and operational dimensions of cyberattack and cyberexploitation are quite different, and this report distinguishes between these two.

Cyberattack has a variety of military applications (discussed in Chapter 3) and may be useful for covert action (discussed in Chapter 4). In addition, cyberattack is conceivably a tool that law enforcement agencies or even the private sector might wish to use under some circumstances (discussed in Chapter 5).

As suggested in the previous section, cyberattack sometimes arises in the context of defending U.S. computer systems and networks. Passive defensive measures such as hardening systems against attack, facilitating recovery in the event of a successful attack, making security more usable and ubiquitous, and educating users to behave properly in a threat environment are important elements of a strong defensive posture.⁷ Nevertheless, for the defense to be successful, these measures must succeed every time the adversary attacks. The adversary's attack need succeed only once, and an adversary that pays no penalty for failed attacks can continue attacking until he or she succeeds or chooses to stop. This places a heavy and asymmetric burden on a defensive posture that employs only passive defense.

If passive defense is insufficient to ensure security, what other approaches might help to strengthen one's defensive posture? One possibility is to eliminate or degrade an adversary's ability to successfully prosecute an attack. In that case, the attack is ultimately less successful than it might otherwise have been because the defender has been able to neutralize the attack in progress (or perhaps even before it was launched).

A second possibility is to impose other costs on the adversary, and such a strategy is based on two premises. First, the imposition of these costs on an attacker reduces the attacker's willingness and/or ability to initiate or to continue an attack. Second, knowledge that an attack is

⁷ The broad topic of steps that might be taken to improve passive cyberdefenses and to enhance resilience of U.S. computer systems and networks is not part of this report. There are many important technology and policy issues in the domain of cyberdefense, but many other works have addressed these issues. For a sampling of relevant National Research Council reports on this topic, see Footnotes 1 and 2 in the Preface to this report. Other important reports include President's Information Technology Advisory Committee, *Cyber Security: A Crisis of Prioritization*, National Coordination Office for Information Technology Research and Development, Washington, D.C., February 2005; and Commission on Cybersecurity for the 44th Presidency, *Securing Cyberspace for the 44th Presidency*, Center for Strategic and International Studies, Washington, D.C., 2008.

BOX 1.2 Terminology Related to Cyberattack¹

A wide variety of terms in the literature have definitions that overlap with the definitions used in this report. (It is perhaps emblematic of the state of discussion today that there is no standard and widely accepted term that denotes attacks on computer systems and networks.) For example:

- The term “information operations” was used by the Joint Chiefs of Staff in 1998 to denote “actions taken to affect adversary information and information systems while defending one’s own information and information systems.” Information operations were characterized as offensive or defensive, where “offensive information operations” were conducted to affect adversary decision makers and achieve or promote specific objectives. The JCS used the term “information warfare” to refer to information operations conducted during time of crisis or conflict (including war) to achieve or promote specific objectives over a specific adversary or adversaries.²

- The term “network attack” is used by the U.S. Air Force Materiel Command’s Electronic Systems Center to refer to “the employment of network based capabilities to destroy, disrupt, corrupt, or usurp information resident in or transiting through networks.”³

- The term “offensive information warfare” was used by Dorothy Denning to describe an operation that “targets or exploits a particular information resource with the objective of increasing its value to the offensive player and decreasing its value to the defensive player.”⁴

- The term “information warfare” has been used often, but with a variety of meanings.⁵ For example, the term is used by the Center for Strategic and International Studies to denote data attack, such as propaganda, disinformation, data overload, and spam; software attack using computer viruses, Trojan horses, or trapdoors; hacking, i.e., penetration, unauthorized use, and/or snooping in other computer systems; and physical kinetic or directed energy attacks against information systems.⁶ By contrast, Ryan and Ryan define information warfare as “the application of destructive force on a large scale against information assets and systems, against computers and networks which support the air traffic control systems, stock transactions, financial records, currency exchanges, Internet communications, telephone switching, credit record, credit card transactions, the space program, the railroad system, the hospital systems that monitor patients and dispense drugs, manufacturing process control systems, newspaper and publishing, the insurance industry, power distribution and utilities, all of which rely heavily on computers.”⁷ Ryan and Ryan also note that “Information warfare is, first and foremost, warfare. It is not information terrorism, computer crime, hacking or commercial or state sponsored espionage using networks for access to desirable information.”

- The term “information attack” is used by Davis Brown, a former deputy judge advocate for the U.S. Defense Information Systems Agency, to focus on information or information systems as the object, means, or medium of attack.⁸

- The terms “offensive cyber operations” and “offensive cyberspace operations” are sometimes heard in discussions with military officials and are apparently used to denote one or more actions, perhaps taken over a period of time, to disrupt, deny, degrade, or destroy information resident in computers and computer networks, or the computers and networks themselves.⁹ Offensive cyber or cyberspace operations apparently extend beyond computer network attack (for example, they include computer network exploitation) and recognize the possibility that an extended offensive campaign might be waged in cyberspace involving multiple cyberattacks.

- The term “computer network attack” was adopted by the Joint Chiefs of Staff in 2006 to refer to “actions taken through the use of computer networks to disrupt, deny, degrade, or destroy information resident in computers and computer networks, or the computers and networks themselves.”¹⁰ In 2006, the Joint Chiefs of Staff also eliminated the term “information warfare” and the distinction between “offensive” and “defensive” information operations.

After considering the plethora of terms used in this domain, the committee settled on “cyberattack” as the term best describing the primary focus of this report.

¹ This description of the various terms is derived in part from Davis Brown, “A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict,” *Harvard International Law Journal*, 47(1):179-221, Winter 2006.

² Joint Chiefs of Staff, Joint Publication No. 3-13, Joint Doctrine for Information Operations, Oct. 9, 1998.

³ See Broad Agency Announcement (BAA ESC 07-0001) on *Network Warfare Operations Capabilities (NWOC): Technology Concept Demonstrations*, May 31, 2007.

⁴ Dorothy E. Denning, *Information Warfare and Security*, Addison-Wesley Longman Ltd., Essex, UK, 1999.

⁵ For a review of such definitions, see Chapter 1 of Gregory Rattray, *Strategic Warfare in Cyberspace*, MIT Press, Cambridge, Mass., 2001.

⁶ *Cybercrime Cyberterrorism Cyberwarfare: Averting an Electronic Waterloo*, Center for Strategic and International Studies, 1998.

⁷ Daniel and Julie Ryan, “Protecting the NII against Infowar,” in Winn Schwartz, *Information Warfare*, Thunder’s Mouth Press, 1996.

⁸ Davis Brown, “A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict,” *Harvard International Law Journal* 47(1):179-221, Winter 2006.

⁹ For example, the U.S. Air Force Cyber Command writes that “Cyberspace favors offensive operations. These operations will deny, degrade, disrupt, destroy, or deceive an adversary. Cyberspace offensive operations ensure friendly freedom of action in cyberspace while denying that same freedom to our adversaries. . . . As an adversary becomes more dependent on cyberspace, cyberspace offensive operations have the potential to produce greater effects.” See Air Force Cyber Command Strategic Vision, undated document (probably 3 March 2008), available at <http://www.afcyber.af.mil/shared/media/document/AFD-080303-054.pdf>.

¹⁰ Joint Chiefs of Staff, Joint Publication No. 3-13, Joint Doctrine for Information Operations, February 13, 2006.

costly to an attacker deters other parties from attempting to attack—and advance knowledge of such a possibility may deter the original adversary from attacking in the first place. There are in general many options for imposing costs on an adversary, including economic penalties such as sanctions, diplomatic penalties such as breaking of diplomatic relations, and even kinetic military actions such as cruise missile strikes. In-kind military action—a counter-cyberattack—is also a possibility.

Both of these possible actions—neutralization of an attacker’s ability to attack and the imposition of costs on the attacker for the attack—are often captured under the rubric of active defense. But actions taken in the name of active defense might well be seen as offensive acts.

Consider the act of Zendia⁸ probing a computer system or network belonging to Ruritania to gather information about it (what ports are open, what services are protected or available for use, the IP addresses of various machines on it, what operating systems are in use, and so on). If Zendia has *already* been the target of a cyberattack launched from Ruritania, Zendia may plausibly regard its probes of computer systems in Ruritania as part of a *defensive* reaction to the attack—gathering information about the systems involved in an attack may be important for characterizing its scale and intent. But Ruritania may regard such a probe as a hostile action by Zendia against it, because such probes can be used to develop information useful in a cyberattack.

The inadequacy of passive defense suggests that the national debate over cybersecurity necessarily includes a consideration of attack options for defensive purposes. Furthermore, once an attack capability is required to conduct active cyberdefense, and once a nation has the capability for active defense, it is also possible for that nation to use an attack capability for other, non-defensive purposes. Attack capabilities may under some circumstances also contribute to deterrence—a relationship that is explicated in more detail in Chapter 9.

Given the possibility that cyberattack capabilities might be useful to the U.S. government for many purposes (including active defense), a host of policy issues arise that do not arise if passive defense is the only defensive option under consideration. Box 1.3 provides an analogy to describe how policy issues inevitably emerge from any government consideration of offensive options.

⁸ Note to the reader: When the name of a nation is needed in this report, the names “Zendia” and “Ruritania” are used as stand-ins. Depending on context, these nations may be a near-peer nation-state with military and economic stature and power comparable to that of the United States; a small, relatively undeveloped nation; or something in between. Generally in this report, Zendia is an adversary of the United States.

BOX 1.3 Policy Issues That Flow from Government Use of Guns

In order for society to defend itself against armed criminals, one policy choice would be to focus on passive defense against guns—bulletproof vests might be distributed to the populace. Criminals might then invest in more powerful guns that could shoot through bulletproof vests. In response, the government might then support research into techniques for developing stronger, more difficult-to-penetrate armor or initiate programs to provide bulletproof vests to more citizens more quickly and educate them about how to use bulletproof vests properly.

Such policy responses are much simpler than those arising from a situation in which police are themselves armed. Governments that arm police officers must be concerned about:

- *Training.* Police officers must have a level of training and expertise in the use of firearms adequate for most situations they will encounter in their day-to-day work.
- *Rules of engagement.* Police officers must follow pre-established rules of engagement that provide guidance on when the use of firearms is and is not appropriate.
- *Command and control.* Police officers are subject to a chain of command that can grant or withhold permission to discharge firearms.
- *Identification friend-or-foe (IFF),* the process by which police officers determine who or what counts as a legitimate target for their weapons. Because undercover police and criminals often choose to look like ordinary citizens (as a rule, they do not wear distinguishing uniforms), police must exercise great care in determining their targets.
- *Liability.* Police (individual officers and the department itself) may be found liable for civil damages or even subject to criminal penalties if a shooting takes place improperly, and especially if someone is injured or killed by such a shooting.

Note that the fact of police officers carrying guns serves a defensive purpose—protecting the citizenry—even though guns themselves are arguably an offensive weapons technology, i.e., a weapons technology that is designed to inflict harm or damage to a target. The committee makes this gun-related analogy not to address any particular policy issue related to private or criminal or even police usage of guns, but to point out that policy and legal issues inevitably flow from the use of offensive weapons by “good guys.”

1.3 CYBERATTACK IN THE CONTEXT OF AN INFORMATION STRATEGY FOR THE UNITED STATES

U.S. military forces have made great progress in developing and implementing plans for joint integrated operations in the conventional

military sphere, but in the information domain, U.S. doctrine and approaches have left many niches and gaps for adversaries to exploit. The lack of an integrated approach to the information domain has meant that the United States lacks timeliness and synergy in its planning and operations. An integrated approach would spread information and ideas that support U.S. interests and would degrade and disrupt information and ideas abroad that are adverse to U.S. interests (e.g., websites for terrorist recruiting).

Cyberattack is only one dimension of information operations. In practice, many cyberattacks are likely to take place within a large, diverse, and organically interconnected domain in which deception, espionage, covert influence, diversion, interception and disruption of communications, and other information operations will also take place (as discussed in Box 3.3 in Chapter 3). All of these operations can be used in an intertwined and integrated fashion. Espionage can be a precursor to a denial-of-service attack, while denial of service can be used to facilitate espionage by forcing one's adversary to use an insecure mode of communication. And information operations are themselves only one aspect of what might be called an information strategy for pursuing U.S. strategic and security interests.

Advocates of such an information strategy argue that the nature of warfare and conflict is changing, and that information will be central to national security affairs in the future. This argument is based in part on the idea that adversaries—unable to compete with the United States in traditional military domains—will seek to exploit U.S. weaknesses asymmetrically, and that the information domain is one of the most important.

Information is central for two reasons. First, modern societies are based largely on the effective use of large amounts of information—a fact reflected in the increasing ubiquity of and dependence on information technology throughout these societies. Second, the “hearts and minds” of much of the world’s population will be won or lost through the influence gained by appropriately targeted ideas and information. The first point suggests that the information assets (and supporting technologies) of modern societies are a possible point of leverage for adversaries that are less dependent on information. The second point suggests that a predominantly military approach to national security is too narrow, and that the United States would be well served by a much broader strategy that puts hearts, minds, and ideas at its center.

In this view, the United States must integrate strategic/tactical influence and messaging and perception management with a broad spectrum of capabilities for information attack and defense. At the highest level of strategic perspective, the goal of information attack is to get into the mind of the adversary and influence its decision making at critical times and

at all levels. This would include making adversaries question their plans, direction, capabilities, actions, likelihood of success, control, and whether generally they trust their information and knowledge base. At the tactical and operational level, information attack entails destroying, denying, degrading, disrupting, influencing, and corrupting an adversary's ability to see, know, understand, decide, and take action. The goal of information defense is to protect our ability to see, know, understand, decide, and take action.

A coordinated information strategy would integrate a variety of disciplines and specialties, most of which are not integrated today. These include strategic communications, influence, and messaging; public diplomacy; perception management; computer network operations (attack, defense, and exploitation); space control; electronic reconnaissance/warfare; psychological operations; strategic and departmental deception; propaganda, information assurance and infrastructure protection, and counter denial and deception; public affairs; counterintelligence; HUMINT (human intelligence) and OSINT (open source intelligence) activities; imagery and mapping operations; data and information mining; and special operations forces.

1.4 IMPORTANT CHARACTERISTICS OF CYBERATTACK AND CYBEREXPLOITATION

As noted above, cyberattack refers to actions—perhaps taken over an extended period of time—to alter, disrupt, deceive, degrade, or destroy adversary computer systems or networks or the information and/or programs resident in or transiting these systems or networks. Several characteristics of weapons for cyberattack are worthy of note.

- *The indirect effects of weapons for cyberattack are almost always more consequential than the direct effects of the attack.* (Direct or immediate effects are effects on the computer system or network attacked. Indirect or follow-on effects—which may be the primary purpose of a cyberattack—are effects on the systems and/or devices that the attacked computer system or network controls or interacts with, or on the people who use or rely on the attacked computer system or network.) That is, the computer or network attacked is much less relevant than the systems controlled by the targeted computer or network (e.g., a cyberattack that affects a computer controlling an electric power generator will also, and more importantly, affect the generator itself) or the decision making that depends on the information contained in or processed by the targeted computer or network, and indeed the indirect effect is often the primary purpose of the attack. Thus, the scale of damage of any given cyberattack can range from

the trivial to the enormous, depending on the systems and/or information connected to or associated with the target.

- *The outcomes of a cyberattack are often highly uncertain.* Minute details of configuration can affect the outcome of a cyberattack, and cascading effects often cannot be reliably predicted. One consequence can be that collateral damage and damage assessment of a cyberattack may be very difficult to estimate.

- *Cyberattacks are often very complex to plan and execute.* Cyberattacks can involve a much larger range of options than most traditional military operations, and because they are fundamentally about an attack's secondary and tertiary effects, there are many more possible outcome paths whose analysis often requires highly specialized knowledge. The time scales on which cyberattacks operate can range from tenths of a second to years, and the spatial scales may be anywhere from "concentrated in a facility next door" to globally dispersed.

- *Compared to traditional military operations, cyberattacks are relatively inexpensive.* The underlying technology for carrying out many types of cyberattacks is widely available, inexpensive, and easy to obtain. An attacker can compromise computers belonging to otherwise uninvolved parties to take part in an attack activity; use automation to increase the amount of damage that can be done per person attacking, increase the speed at which the damage is done, and decrease the required knowledge and skill level of the operator of the system; and even steal the financial assets of an adversary to use for its own ends. On the other hand, some cyberattack weapons are usable only once or a few times.

- *The identity of the originating party behind a significant cyberattack can be concealed with relative ease, compared to that of a significant kinetic attack.* Cyberattacks are very difficult to attribute to any particular actor and are thus easy to conduct with plausible deniability—indeed, most cyberattacks are inherently deniable. Cyberattacks are thus also well suited for being instruments of catalytic conflict—instigating conflict between two other parties.

Many of the operational considerations for cyberexploitation are similar to those for cyberattack. Like cyberattack, a successful cyberexploitation requires a vulnerability, access to that vulnerability, and a payload to be executed—the only difference is in the payload to be executed. These similarities often mean that a targeted party may not be able to distinguish easily between a cyberexploitation and a cyberattack—a fact that may result in that party's making incorrect or misinformed decisions. The primary technical requirement of a cyberexploitation is that the delivery and execution of its payload be accomplished quietly and undetectably. Secrecy is often far less important when cyberattack is the

mission, because in many cases the effects of the attack will be immediately apparent to the target.

1.5 ILLUSTRATIVE APPLICATIONS OF CYBERATTACK

Cyberattack can be used to support many traditional military operations, such as the disruption of adversary command, control, and communications; suppression of adversary air defenses; degradation of adversary smart munitions and platforms; and attack of adversary warfighting or warmaking infrastructure (the adversary defense industrial base). Cyberattack might be used to augment a kinetic attack or to enable it to succeed, or to defend a friendly computer system or network by neutralizing the source of a cyberattack conducted against it. Cyberattack could also be used to achieve military deception. For example, by assuming control of a computer used by a senior intelligence analyst, a cyberattack could send bogus e-mail traffic to that analyst's clients. The contents of these e-mails could easily provide misinformation regarding the military capabilities, intentions, locations, and operations of friendly forces.

From a strictly technical perspective, cyberattack has several attributes that are well suited for the shadowy world of intelligence. For example, as noted above, attribution of a cyberattack is usually quite difficult. The effects of a cyberattack may not become visible to the victim for long periods of time, if ever. And the range of possible options is very large, so that cyberattack-based operations might be set in motion to influence an election, instigate conflict between political factions, harass disfavored leaders or entities, or divert money. Such operations can fall into the category of covert action, which by law is defined as political, economic, propaganda, or paramilitary activities and is usually designed to influence governments, events, organizations, or persons in support of foreign policy in a manner that is not necessarily attributable to the U.S. government.

1.6 THE LEGAL FRAMEWORK GOVERNING CYBERATTACK

The committee's view of the basic framework for the legal analysis of cyberattack is based on the principle that notions related to "use of force" and "armed attack" (terms of special relevance to the Charter of the United Nations) should be judged primarily by the effects of an action rather than its modality. That is, the fact that an attack is carried out through the use of cyberweapons rather than kinetic weapons is far less significant than the effects that result from such use, where "effects" are understood to include both direct and indirect effects.

Accordingly, cyberattack should be judged according to the principles

of the law of armed conflict (LOAC) and the UN Charter, encompassing both *jus ad bellum* (law governing the legality of going to war) and *jus in bello* (law governing behavior during war) with the understanding that new analytical work is needed to understand how these principles do or should apply to cyberweapons. For example, some of the more problematic cases involving cyberattack include the following:

- *Conflicts that do not fall under the presumption of nation-to-nation conflict between national military forces.* When the law of armed conflict was first articulated, only nation-states had the ability to wage war. Because cyberattack weapons are inexpensive and easily available, non-state actors (e.g., terrorist groups, organized crime) are capable of engaging in armed conflict through the use of cyberweapons, as are individuals acting on their own with putatively “patriotic” motivations or with criminal intentions. Even in non-government hands, these weapons include some that are as capable of doing great harm as those available to governments. Thus, the lines between state, non-state, and individual attackers are unclear in a legal regime that distinguishes between LOAC on the one hand and national criminal laws on the other.

- *The exception for espionage.* The LOAC presumes that a clear distinction can be drawn between the use of force and espionage, where espionage is avowedly not a use of force. However, the distinction between cyberattack and cyberexploitation may be very hard to draw from a technical standpoint, and may lie primarily in the intent of the user.

- *The emphasis on notions of territorial integrity.* A target in cyberspace may be known only through an electronic identifier, such as an IP address or a MAC address. (A component’s Media Access Control address—generally known as a MAC address—is a quasi-unique identifier assigned to most network adapters or network interface cards (NICs) by their manufacturer for identification.) To what extent should the physical location of a computer matter in determining whether it is a legitimate military target that may be subject to cyberattack? Also, the effects of attacking a given computer may not be felt at all in the immediate geographic vicinity of the computer, thus raising the question of which geographic location is relevant to the determination of legitimacy for attack.

1.7 THE DYNAMICS OF CYBERCONFLICT

The escalatory dynamics of armed conflict are thought to be understood as the result of many years of thinking about the subject. The dynamics of cyberwarfare are less well understood. This report speculates on some of the factors that might influence the evolution of a cyberconflict.

For major nation-states with significant kinetic and cyber capabilities at their disposal, some of the important questions to be addressed include the following:

- *Crisis stability.* What is the analog of crisis stability in cyberconflict? What are the incentives for preemptive cyberattack? Crisis stability refers to the condition in which even in a crisis, neither side has an incentive to escalate the conflict.

- *Resolving the tension between a policy need for rapid response and the technical reality that attribution of a cyber action is a time-consuming task.* Shortening the time for investigation may well increase the likelihood of errors being made in a response (e.g., responding against the wrong machine, launching a response that has large unintended effects).

- *Preventing cyberconflict from escalating to physical space.* Given that cyberattacks are likely to occur in the early stages of a conflict, how can cyberconflict between nations be limited to conflict in cyberspace? How should cyberattack be scoped and targeted so that it does not lead an adversary to escalate a conflict into kinetic conflict? How can a modestly scoped cyberattack conducted by a government be differentiated from the background cyberattacks that are going on all of the time?

- *The complicating presence of non-state actors.* How can “freelance” activities on the part of “patriotic hackers” be minimized or curtailed?

- *Termination of cyberconflict.* How would two nations engaged in cyberconflict indicate that they have ceased cyberattacks against each other?

- *The role of transparency.* What is the role of transparency in promoting crisis stability and conflict limitation in cyberspace?

- *Catalytic cyberconflict.* How can catalytic cyberconflict be avoided? (Catalytic conflict refers to the instigation of conflict between two parties at the behest or initiative of a third party.)

For non-state actors such as terrorist or criminal groups, two primary issues relate to identification of the appropriate party against which to retaliate, and the availability of cyber and/or kinetic targets whose destruction might cause pain or meaningful damage to the terrorist or criminal group. At the same time, nations hosting such groups might have plausible targets, and the assistance of those nations in acting against such groups might be obtained.

1.8 FINDINGS

This section presents the committee's findings and recommendations, along with supporting arguments that summarize material contained in later chapters.

1.8.1 Technologies as Instruments of U.S. National Policy

Once a need has been established, a sound stance regarding the use of most technologies as an instrument of U.S. national policy rests on four pillars:

- *Capabilities to use the technology in a variety of situations and contexts.* That is, the technology must be sufficiently well developed and robust to be usable in ways that advance U.S. national interests. When new technologies are in their infancy, unproven extravagant claims are often made about their putative effectiveness—but in some cases, such claims do turn out to be valid.

- *Policy guidance for when and how such capabilities should be exercised.* Policy guidance can be expressed in many forms, including statutory and/or common law, regulations and directives, ethical standards, acquisition decisions, military doctrine, and so on. As a general rule, these different expressions of national policy should reinforce, or at least be consistent with, each other. But since the U.S. government, like all governments, has multiple loci for policy formation, such consistency is not always found in national policy.

- *Decision-making mechanisms for implementing policy guidance in an operational sense regarding the actual use of the capabilities available.* For example, when a crisis occurs, a well-organized government will have clear and transparent mechanisms in place for directing that various actions be taken in response. One central element of such mechanisms is necessarily focused on reconciling competing interests and equities that may be present among the various stakeholders represented in the government and/or in the private sector.

- *Oversight to ensure consistency between actual use and policy guidance.* In large bureaucracies, maintaining consistency between policy and actual practice or use is often difficult, and oversight is necessary to ensure such consistency. In practice, oversight closes the feedback loop between outcomes and policy, and provides indicators of whether a policy is working to advance national interests.

The first of these elements is inherent in the technology. But the remaining three elements emerge only from the organizations and people who must determine how any given technology is to be used—and such

emergence almost always trails the development of technological capabilities. Furthermore, a rapid pace of change in technologies relevant to warfare almost always changes the nature of warfare itself, and thus military doctrine and concepts of use must also adjust to the realities of new technologies.

Regarding cyberattack, consider that the World Wide Web was invented in the early 1990s⁹ and personal computers went mainstream for the general public less than 30 years ago with the introduction of the IBM PC in 1981. Over a billion people use cell phones today,¹⁰ and wireless services are growing exponentially. Accurate location and velocity information for vehicles is more available than ever before through GPS and similar systems. Taking into account the speed at which organizations such as national governments change, it is not surprising that policy guidance, decision-making mechanisms for operational use, and oversight mechanisms for cyberattack have not been fully developed. These elements are the primary focus of the findings and recommendations that follow.

As for today's policy context, policy and guidance are evolving rapidly at the time of this writing (early 2009). A number of reports have been recently issued speaking to the importance of cybersecurity to the nation. These reports have either obliquely or explicitly referred to the importance of integrating defensive activities with offensive activities in cyberspace. The outgoing administration launched the \$40 billion Comprehensive National Cybersecurity Initiative (CNCI), which reportedly takes seriously this notion of integration. The organization of the Department of Defense for cyber operations is in flux as well, as different agencies and services make their cases for significant roles regarding the attack mission.

1.8.2 Overarching Findings

Finding 1: The policy and organizational issues raised by U.S. acquisition and use of cyberattack are significant across a broad range of conflict scenarios, from small skirmishes with minor actors on the international stage to all-out conflicts with adversaries capable of employing weapons of mass destruction.

⁹ See <http://groups.google.com/group/alt.hypertext/msg/395f282a67a1916c>.

¹⁰ More precisely, the number of mobile telephone subscriptions globally reached 3.3 billion in November 2007. See Reuters, "Global Cellphone Penetration Reaches 50 Pct," November 29, 2007, available at <http://investing.reuters.co.uk/news/articleinvesting.aspx?type=media&storyID=nL29172095>.

The statement above represents the primary finding of the committee. All of the findings in this section are elaborations of this primary finding.

While the immediate effects of cyberattack are unlikely to be comparable to the effects of weapons of mass destruction (for example, nuclear, chemical, or biological weapons), a large-scale cyberattack could massively affect the functioning of a society and lead to many indirect casualties. Conversely, it is possible to imagine that certain cyberattacks might be executed on a smaller scale and with a lower degree of lethality than might be expected if kinetic weapons were used for equivalent military purposes. Thus the policy implications of cyberattack have certain commonalities across the range from non-lethal engagements to wars involving the use of weapons of mass destruction.

To the extent that new technologies afford new capabilities, they imply new policy challenges about how to develop, acquire, and use them; about who should use them and who should decide about using them; and indeed about how to think about them. But the policy issues associated with cyberattack are of particular urgency today, because the amount and degree of conceptualization and understanding in the policy-making community about these issues relative to their potential significance is much lower than is the case with almost any other weapon in the U.S. arsenal. In other words, the state of policy formation regarding cyberattack is still in its infancy compared to policy regarding most other weapons, even though the availability and proliferation of cyberattack technologies is a technological watershed. And it is the committee's belief that the issues surrounding cyberattack extend far beyond the traditional responsibilities of the Department of Defense and the intelligence community and touch national interests such as diplomacy and foreign relations, law enforcement, and commerce and trade.

Finally, the committee notes that the goals of a cyberattack (i.e., the alteration, disruption, deception, degradation, or destruction of a computer system or network) may sometimes be accomplished by more traditional kinetic means. Any planner contemplating the destruction of an adversary computer or network would have to think about both cyberattack and kinetic attack options. But there is a well-developed body of doctrine and guidance regarding kinetic options, and so the committee has not specifically examined or presented the kinetic perspective in any systematic way in this report.

Finding 2: The availability of cyberattack technologies for national purposes greatly expands the range of options available to U.S. policy makers as well as to policy makers of other nations.

Cyberattack technologies can have a broad range of effects and impacts (Chapter 2). They are thus quite flexible, and for example can sometimes be operated reversibly or irreversibly and in a lethal/destructive or non-lethal/non-destructive manner depending on the specific technology involved. In addition, cyberattack technologies have a largely clandestine character and are relatively inexpensive. These characteristics—flexibility, clandestine nature, and low cost—can be helpful in many applications by the military, intelligence, and law enforcement communities.

An important consequence of the broad range of possible effects and impacts is that cyberattack as an instrument of national policy has both offensive and defensive implications (as noted in Section 1.2 and Chapter 2), and both tactical and strategic implications as well (Chapter 9). Furthermore, much of the supporting technology for characterizing ongoing cyberattacks (e.g., detection, warning, attack assessment) is relevant to both offense and defense, and to tactical and strategic planning and decision making. Such dualities can be found for kinetic technologies as well, but they are front and center in thinking about cyberconflict.

The fact that cyberattack technologies can have a broad range of effects and impacts also means that their use may sometimes result in unanticipated, unforeseen, or unintended consequences. Concerns about these unanticipated consequences may (and perhaps should) inhibit the use of cyberattack under some circumstances. Finding 18 notes the uncertainties associated with the effects of many kinds of cyberattack.

In addition, the nature of cyberattack technologies is that they are available to other nations and non-state actors as well as to the United States—a fact that results in the plentitude of vulnerabilities to the U.S. critical infrastructure and U.S. military documented in so many reports.¹¹

Finding 3: Today’s policy and legal framework for guiding and regulating the U.S. use of cyberattack is ill-formed, undeveloped, and highly uncertain.

To date, national policy regarding cyberconflict has focused mostly on the defense of friendly computer systems and networks against cyberattack, although by most accounts the information technology infrastructure of the United States is still quite vulnerable and policy for cyberdefense

¹¹ See, for example, President’s Information Technology Advisory Committee, *Cyber Security: A Crisis of Prioritization*, National Coordination Office for Information Technology Research and Development, Washington, D.C., February 2005.

is still uncertain.¹² But the United States has no comprehensive publicly stated strategic national policy apart from a criminal framework concerning how it will regard cyberattacks conducted against the United States or how it might use cyberattack in support of U.S. interests. The most relevant international legal framework—the law of armed conflict combined with the Charter of the United Nations—was formulated in an era that long predates the information age and cyberattack, and although the principles of the LOAC framework still apply (Finding 6), the specifics of applying the principles to cyberattack are sometimes uncertain.

These points illustrate the lack of a shared conceptual understanding about the full spectrum of issues regarding cyberattack among all of the stakeholders—military, intelligence, law enforcement authorities, and the private sector. Such a shared understanding is a prerequisite for responsible decision making about this topic.

The undeveloped and uncertain nature of this legal and policy framework poses a number of dangers for the United States, not the least of which is that policy and law developed in a time of (or in response to) crisis are often—some might argue usually—hastily formulated and thus incompletely considered. Crisis may also bias the policy consideration in undesirable ways. For example, arguments in favor of a particular course of action can be artificially bolstered by crisis, and arguments against that course of action artificially suppressed, thus bypassing the weighing of tradeoffs that characterizes non-crisis decision making. And unsound policy formulated and implemented during crisis may prove difficult to change or reverse when the crisis has passed.

Finding 4: Secrecy has impeded widespread understanding and debate about the nature and implications of U.S. cyberattack.

The relatively recent emergence of cyberattack technologies and the resulting dearth of associated policy, law, and ethics raise some very important issues for all sectors of society. Nevertheless, a full public discussion of these issues has yet to coalesce, and classification of such topics as being at secret or higher levels has left U.S. government thinking on these issues highly opaque. Such opacity has many undesirable consequences:

- Neither the potential importance and usefulness of cyberattack as

¹²See, for example, National Research Council, *Information Technology for Counterterrorism: Immediate Actions and Future Possibilities*, The National Academies Press, Washington, D.C., 2003, and National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington D.C., 2007.

an instrument of national policy nor the potential perils and pitfalls of using cyberattack are well understood outside niches in the defense and intelligence communities. Secrecy about policy relevant to cyberattack inhibits public scrutiny and thus increases the likelihood that policy will be formulated with narrow parochial or short-term interests foremost in mind.

- Programs to develop cyberattack capabilities are classified and dispersed throughout many program elements within the Department of Defense, with the result that overall capabilities may not be widely known even among those with the necessary clearances. Effective congressional oversight that goes beyond a few individuals on the relevant committees is also inhibited.

- Unclassified programs to develop stronger or more effective defensive capabilities do not benefit from the insights derived from knowledge of cyberattack. Yet it is well known that many intellectual and programmatic synergies are possible when experts in defense and attack can collaborate.

- Independent research and investigation about the topic is inhibited, in particular for two groups: non-military/non-government researchers and DOD/intelligence community personnel (both uniformed and non-uniformed) who do not now but may in the future need to know about this area.

—For the first group, the loss of independent non-governmental analysis increases the likelihood that the full array of national and international intellectual capital will not be brought to bear on the issue, thereby depriving policy makers of its potential contributions to understanding the issue. (In this regard, the committee notes a 50-year history of independent non-government analysis that has made important contributions to the formulation of U.S. policy regarding nuclear, chemical, and biological weapons.)

—For the second group, it is not reasonable to expect individuals placed into responsible positions to get up to speed quickly if they do not have the basic and fundamental background knowledge needed. Yet this is precisely what is implied by the current regime of secrecy surrounding cyberattack—DOD/intelligence community personnel in other assignments have no reasonable opportunity to be exposed to the basic policy issues involved in cyberattack (because they have no “need to know” in their current duty assignments), and they are expected to be in a position to make sound policy judgments when they fill their cyberattack billets. Moreover, personnel in non-cyberattack assignments need to comprehend the basic policy issues involved in cyberattack so that they are able to

understand and assess what these policy issues for cyberattack mean for the responsibilities of their current positions.

- Dissemination even of unclassified information is inhibited by the broad classification of the cyberattack topic. Guidelines for the protection of classified information force individuals with clearances and access to such information to be certain that information they discuss publicly is indeed unclassified—otherwise, they are obligated to treat any material received in classified settings as classified even if the material in question is in fact unclassified. The result is that such individuals are reluctant to talk publicly about such issues at all.
- Professional military education cannot explore the cyberattack topic in any meaningful sense. Nevertheless, professional military education is one of the most important venues in which those military personnel unfamiliar with critical topics can learn about them.
- Secrecy has also inhibited discussion of issues related to cyberattack outside the defense context. For secrecy as well as other reasons, discussion about the pros and cons of cyberattack as a component of a comprehensive defense has been inhibited and delegitimized. Greater public discussion of cyberattack in a military context is likely to spur greater discussion of related issues in non-military contexts.

Finding 5: The consequences of a cyberattack may be both direct and indirect, and in some cases of interest, the indirect consequences of a cyberattack can far outweigh the direct consequences.

To the extent that there has been any public discussion about cyberattack, the full range of possible effects and consequences of cyberattack is often not addressed. In fact, cyberattacks can have a very broad range of consequence, from barely noticeable by careful observers to immediately significant on a global scale. For this reason, any discussion of cyberattack in use must address its effects. Furthermore, since the information available to attackers may well be limited, there will also be some range of uncertainty about the extent and nature of a cyberattack's effects.

A full consideration of a cyberattack's effects necessarily includes both direct (immediate) and indirect (follow-on) effects (Section 1.4). Direct or immediate effects are effects on the computer system or network attacked. Indirect or follow-on effects are effects on the systems and/or devices that the attacked computer system or network controls or interacts with, or on the people that use or rely on the attacked computer system or network.

Another dimension of the effects issue relates to the time scale on which the effects of a cyberattack will be manifest. Some cyberattacks

target data or software, and such attacks are usually more easily and more rapidly reversed, at least in part, than are kinetic attacks on objects that are energetically disassembled (blown up). For example, if backup media are easily accessible, it may be possible to restore corrupted data and software to their pre-attack states relatively quickly and with only minimal losses. An attacker might have the ability to restore data as well (e.g., some cyberattacks call for the in-place encryption of data and decrypting it only after hostilities terminate). But the indirect or follow-on effects are generally not as easily reversible in much the same way that the direct effects of kinetic attacks are generally not easily reversible. A cyberattack on the computer controlling a generator or a dam that was in fact intended to disable the generator or the dam could cause the generator to burn itself out or the dam to release its floodgates too soon. Such effects are kinetic in nature, and thus are only as reversible as their underlying physical structures are replaceable or repairable. (Even in the case where a cyberattack is intended to confuse the enemy (e.g., by altering data) rather than to cause a kinetic or physical effect, the ultimate results of that confusion are likely to be difficult to reverse.) Depending on the nature of the cyberattack, the extent of reversibility may be an additional (and possibly significant) factor in undertaking any analysis of its effects.

One important consequence of Finding 5 is that policy makers and operational commanders cannot assume that cyberattacks are non-lethal simply because they target computers or networks—a fact that provides further support for Finding 1. The full scope of effects, both direct and indirect, must be taken into account in a determination of the lethality and other consequences of any given cyberattack—and this is true for attacks launched by the United States as well as for attacks directed against the United States.

A second consequence is that not all cyberattacks constitute “cyberwarfare.” As a form of warfare, cyberwarfare automatically brings in all of the associated legal and ethical constructs associated with the term, and they may not apply in all cases of cyberattack. Furthermore, cyberattacks should not be conflated with cyberexploitations, as they often are in the popular press and in lay discussions of the topic (Box 1.4).

1.8.3 Legal and Ethical Findings

Much of today’s current thinking about how to engage in armed conflict originated a century ago, and thus it is not surprising that today’s international law—and especially the law of armed conflict—may not be entirely adequate to handle all of the implications of cyberattack technologies that have emerged only in the last few decades. The same is true of

BOX 1.4 Conflation of Cyberattack and Cyberexploitation— Negative Consequences

Cyberattack and cyberexploitation are often conflated in public discourse, and in particular cyberexploitations are reported and discussed using the term “cyberattack.” For example:

- *Congress.* Representative Frank Wolf (R-VA) stated on the House floor in June 2008 that “In August 2006, four of the computers in my personal office were compromised by an outside source. On these computers was information about all of the work I have done on behalf of political dissidents and human rights activists around the world. . . . The FBI revealed that the outside sources responsible for this *attack* [emphasis added] came from within the People’s Republic of China.”¹
- *News organizations.* A *Time* magazine article of 2005 stated that “Car-penter had never seen hackers work so quickly, with such a sense of purpose. They would commandeer a hidden section of a hard drive, zip up as many files as possible and immediately transmit the data to way stations in South Korea, Hong Kong or Taiwan before sending them to mainland China. They always made a silent escape, wiping their electronic fingerprints clean and leaving behind an almost undetectable beacon allowing them to re-enter the machine at will. An entire *attack* [emphasis added] took 10 to 30 minutes.”²
- *National laboratories.* In December 2007, the Oak Ridge National Laboratory posted a notice labeled *Potential Identity Theft* stating that “Oak Ridge National Laboratory (ORNL) recently experienced a sophisticated *cyber attack* [emphasis added] that appears to be part of a coordinated attempt to gain access to computer networks at numerous laboratories and other institutions across the country. A hacker illegally gained access to ORNL computers by sending staff e-mails that appeared to be official legitimate communications. When the employees opened the attachment or accessed an embedded link, the hacker planted a program on the employees’ computers that enabled the hacker to copy and retrieve information. The original e-mail and first potential corruption occurred on October 29, 2007. We have reason to believe that data was stolen from a database used for visitors to the Laboratory.”³

domestic law—this too has lagged behind the times in coming to terms with the implications of new cyberattack technologies.

Finding 6: The conceptual framework that underpins the UN Charter on the use of force and armed attack and today’s law of armed conflict provides a reasonable starting point for an international legal regime to govern cyberattack. However, those legal constructs fail to account for non-state actors and for the technical characteristics of some cyberattacks.

The committee believes that conflating these terms does not contribute to an informed public discussion of cyberattack or the broader discussion of cybersecurity. Indeed, such conflation has a number of negative consequences:

- It overstates the actual threat, thus inflaming public passion and beating the drums of war unnecessarily. It is certainly true that cyberexploitations are not friendly acts, but they are not armed attacks either. Most nations engage in espionage even against allies and neutral nations without it leading to war or even armed conflict, and cyberexploitation is in essence a form of espionage.
- Calling a cyberexploitation an attack may imply in the public mind an immediate right to counterattack—perhaps through cyber means or perhaps through kinetic means—even though the action in question would not properly be regarded as a military attack. Thus, if policy makers lump together cyberexploitations and real cyberattacks as “cyberattack,” they may well be impelled to counterattack with more force than is appropriate under the circumstances.
- Calling cyberexploitation a cyberattack could prejudice U.S. positions and interests in future cyber arms control talks. With an overly broad definition, the United States might find itself unwilling to ratify a treaty in order to preserve certain capabilities that fall short of actual attack, and thus end up outside international norms even when it might not object to limiting certain attack capabilities.

¹ See <http://wolf.house.gov/?sectionid=211§iontree=7,211&itemid=1213>. The Congressional Record transcript can be found at http://www.fas.org/irp/congress/2008_cr/wolf061108.html.

² By Elaine Shannon, “The Invasion of the Chinese Cyberspies (and the Man Who Tried to Stop Them),” *Time*, August 29, 2005, available at <http://www.time.com/time/magazine/article/0,9171,1098961,00.html>.

³ See <http://www.ornl.gov/identifytheft/>.

The committee believes that the conceptual framework that underpins the UN Charter and today’s law of armed conflict regarding the use of force and armed attack is generally consistent with the notion that the effects of an action rather than the modality of that action are the primary measure in judging its legality under the UN Charter or LOAC.

Prior to an acknowledged armed conflict, the legal status of any military activity is judged by its effects (regardless of the means) according to the criteria of the UN Charter and *jus ad bellum*. Therefore, if the effects (including both direct and indirect effects) to be produced by a cyberattack would, if produced by other means, constitute an armed attack in the

sense of Article 51 of the UN Charter, it should be treated as an armed attack. Similarly, if a cyberattack would have the same effects as certain governmentally initiated coercive/harmful actions that are traditionally and generally not treated as the “use of force” (e.g., economic sanctions, espionage, or certain covert actions), such a cyberattack should also not be regarded as a use of force.

Article 51 acknowledges the conditional right of a nation to engage in the use of armed force for self-defense, including the situation in which the nation is the target of an armed attack, even without Security Council authorization. Thus, the response to a cyberattack by acts constituting use of force is legal, permitted, and proper only if and when—but definitely if and when—the effect of the initial action is equivalent to the effect of an armed attack. If the initial provocation does not rise to the level of being an armed attack, it is not legal to respond with any act that constitutes the use of force, whether cyberattack or otherwise.

The committee also concurs in the judgment of the U.S. armed forces that during acknowledged armed conflict (notably when kinetic and other means are also being used against the same target nation), military use of cyberattack is governed by all the standard LOAC criteria of *jus in bello*—military necessity (and seeking the destruction only of legitimate targets that make a direct contribution to the enemy’s war effort), proportionality (and thus pursuing offensive action only when the military advantage to be gained by the attack outweighs the collateral damage that would ensue), and distinction (restricting combatants and non-combatants to their legitimate roles in return for the different legal protections afforded to them).

At the same time, the framework underpinning existing law is poorly suited to deal with certain aspects of cyberattack. One major complicating factor in the analysis of cyberattacks is that although the law of armed conflict continues to govern the use of cyberattacks by the U.S. armed forces (and in principle, by other nations as well), LOAC is based on a state-to-state framework and thus largely assumes interstate conflict. But today, and especially in cyberspace, non-state actors (e.g., terrorist groups, organized crime) are entirely capable of engaging in armed conflict, as are individuals acting on their own with putatively “patriotic” motivations—and the lines between state, non-state, and individual attackers are unclear in a legal regime that focuses primarily on LOAC on the one hand and national criminal laws on the other. International agreements, such as the Convention on Cybercrime (Section 7.2.4), will help to increase the effectiveness of criminal law in dealing with cyberattacks, but it is likely that some gray area will always exist between LOAC and criminal law when certain kinds of cyberattack occur.

Of course, the notion that a threat might emanate from a non-state

actor is not unique to the domain of cyberattack. Terrorists seeking to inflict kinetic damage often operate from a neutral nation's territory—and indeed, using cyberattack as an instrument against non-state actors brings into play many of the issues that arise in fighting terrorists. For example, self-defense against attacking parties in neutral territory (discussed in Section 7.2.2.4) can easily become relevant to a decision to launch a counter-cyberattack against a cyberattack apparently emerging from a neutral nation—and the structure of the process used in deciding this case would be very similar to the decision-making process used in deciding whether to launch a kinetic attack against terrorists operating from a failed state.

A second complicating factor is related to various technical characteristics of cyberattacks that may be carried over the Internet. Today, the United States is undertaking major efforts to monitor Internet activity for indications of hostile intent. For example, the DOD, under the auspices of the U.S. Strategic Command, monitors attacks on DOD systems. A variety of computer emergency response teams and commercial anti-virus/worm-detection firms also continually monitor Internet network operations for indications of threat warning. These monitoring efforts are likely to provide some degree of “early warning” for impending cyberattacks conducted over the Internet, although that time may be measured in seconds. As such an attack unfolds, its scope and effects may become clearer as well.

However, these efforts at Internet surveillance will not necessarily reveal planning and preparation of the attack, nor intent or even origin. Neither does surveillance reveal aspects of the attack that take place at protocol levels below the monitoring sensors, or above them if concealed. (For example, a cyberattack may be deliberately designed to hide the extent and nature of the damage it causes. In addition, an adversary's battlefield preparation for a cyberattack (e.g., installing easy-to-use back doors) may be done surreptitiously, thus making it difficult for the victim to know the scope and nature of the preparation.) Because LOAC and the UN Charter presume not only nation-states in conflict but also that the specific nation-states involved are known to all, the difficulty of attributing a cyberattack in its early stages to a particular actor, which may be a state or a non-state actor, remains a major challenge to the current legal regime. Thus, the United States may know that it has suffered an “armed attack” or been the target of a “use of force,” but it may take a long time to determine the party or parties responsible.

Finally, because so much of LOAC and the UN Charter is based on the idea that civilian and military assets can be separated, the intermingling and interconnection of military and civilian information technology assets and the importance of a nation's critical infrastructure to both military and

civilian activities will present challenges to today's LOAC/UN Charter regime. Even with the intent to comply with LOAC and the UN Charter, policy makers will face many difficulties in arriving at sound judgments regarding events involving cyberattack, whether the United States is the victim or the launcher of a cyberattack.

Finding 7: In today's security environment, private parties have few useful alternatives for responding to a severe cyberattack that arrives over a network such as the Internet.

When a private party is the target of a cyberattack that arrives over a network such as the Internet, it has four options for responding. First, it can implement passive measures to strengthen its defensive posture. For example, it can drop functionality on its own systems that the attacker is trying to exploit, reject traffic, and close ports on its firewall. Second, it can report the event to law enforcement authorities, and law enforcement authorities can take appropriate action to try to shut down the cyberattack (e.g., by finding the perpetrator and arresting him, but not by launching a cyber counterattack). Third, it can take self-help measures to further investigate and characterize the source of the cyberattack and then report the information to appropriate law enforcement authorities. Fourth, it can take actions to actively neutralize the incoming cyberattack.

The first two options are generally legal under U.S. domestic law. But the first option may cause the victim to lose the benefit of essential computer and network services and connections. With respect to the second option, law enforcement authorities may not be able to respond effectively on a time scale that will prevent significant immediate harm to the victim, although arrest and prosecution might provide a possible venue for restitution. That is, there appears to be no government agency that has the legal authorization to perform a "harm cessation" function apart from the arrest-and-prosecute mode.

Furthermore, the appropriate and relevant law enforcement authorities are not always easily identified. If a U.S. firm with offices in Japan is cyberattacked in Japan by the Russian mob, are the cognizant law enforcement authorities American (because the firm is a U.S. firm), Japanese (because that was the place where the consequences were manifested), or Russian (because Russia was the national home of the bad attackers)?

If the first two options are not sufficient to keep losses to an acceptable level, the victim might understandably consider the third and fourth options. That is, if the victim is unable to strengthen its defenses without losing essential functionality, and law enforcement authorities cannot prevent further harm, self-help options gain in attractiveness.

However, regarding option three, it may well be illegal under both

the Computer Fraud and Abuse Act (CFAA) and the Electronic Communications Protection Act for the victim to investigate and characterize the attack and attacker by initiating probes of its own, even if such information would be useful for law enforcement authorities in conducting their own investigation.

Option four likely violates the CFAA, which forbids private individuals and organizations to intentionally cause damage in excess of \$5,000, without authorization, to any computer “used in interstate or foreign commerce or communication, including a computer located outside the United States that is used in a manner that affects interstate or foreign commerce or communication of the United States.” Still, limited rights regarding the defense of property intended to prevent continuing harm have traditionally been afforded to victims of attack under some circumstances even in the absence of explicit legislative authority for actions taken in the defense of property. In short, even a private party under continuing cyberattack may itself have some rights to use a cyberattack of its own to stop the incoming cyberattack.

To the best of the committee’s knowledge, defense of property has never been invoked in a defense against charges of violating the CFAA, and so the legal justifiability of such actions is subject to some doubt.

At the same time and regardless of the legality of such actions, exercise of such rights may well be problematic both for the attacked party and for the nation at large from a policy perspective:

- The particular conditions necessary to invoke rights to defense of property are not clearly specified anywhere, and thus some degree of legal uncertainty necessarily attaches to such actions.
- Because many kinds of cyberattack can be transmitted across large distances, cyber actions taken to respond to a cyberattack will almost inevitably invade the premises of the attacker.
- Actions taken by the attacked party may be attributed to the government with responsible authority over it, especially if government standards governing the invocation of such rights are established.
- Given the difficulties of technical attribution of a cyberattacker, a victim undertaking a responsive cyberattack has a non-negligible chance of striking innocent third parties, making defense of property in this context far more problematic than the defense engaged in by a homeowner shooting at a home intruder.

Finding 8: Cyberattack poses challenges to existing ethical and human rights regimes.

As noted in Chapter 7, the laws of armed conflict are based on two

central ethical principles—that the use of force or violence against another state must be justified by “good” reasons for doing so, and that even if violent conflict between nations is inevitable from time to time, unnecessary human suffering should be minimized. To the extent that the laws of armed conflict govern the use of cyberattack, cyberattack is not a *sui generis* phenomenon that is incompatible with these ethical principles.

On the other hand, cyberattack can complicate the application of these ethical principles. For example, the argumentation for Finding 6 noted the complications introduced into today’s legal regime by the dual-use nature of today’s information technology infrastructure. This dual-use nature also complicates ethical judgments that have traditionally been based on the notion of separating civilian and military assets, and the need for making such judgment may well be relevant in situations short of acknowledged armed conflict in which LOAC is held to apply.

The possibility of extended cyberattacks on a society’s information technology infrastructure also raises the question of whether the IT-dependent features of modern society are in any sense essential to life as the citizens of that society know it. For example, the citizens of a large nation often use credit cards to conduct retail transactions. If a cyberattack on the financial infrastructure disrupted the ability of citizens to conduct electronic transactions for an extended period of time without causing large-scale death or destruction of property, what, if any, is the ethical responsibility of the nation launching the attack? Estonia, for example, has gone so far as to declare that Internet access is a fundamental right of its citizenry.¹³ An extended cyberattack campaign against a modern nation that deprived citizens only of such features of modern life (and did not cause large-scale death or destruction of property) might still be reasonably considered a use of force by the attacked nation and the world community and/or a human rights violation of the citizens of the attacked nation by the attacker.

The International Covenant on Civil and Political Rights articulates one current international understanding of human rights. But although a number of its provisions can be argued to be relevant to the cyber domain, it is reasonably clear that the framers of that convention did not take explicit account of the possibility that cyberattacks might affect human rights. The United States has argued that the convention does not apply extraterritorially, and hence it would not regulate U.S. behavior regarding other countries—however, as a practical matter, the role of human rights law during conflict is contested internationally, and there is no reason to expect that cyberconflict will be exempt from this debate.

¹³ Colin Woodward, “Estonia, Where Being Wired Is a Human Right,” *Christian Science Monitor*, July 1, 2003, available at <http://www.csmonitor.com/2003/0701/p07s01-woeu.html>.

Finally, if cyberattack capabilities are seen as providing policy makers with an alternative short of using traditional kinetic armed force in the conduct of their international relations (Section 8.4), they may increase the likelihood that national leaders will choose to intervene when they might otherwise have refrained from intervention. Such an outcome may raise ethical and moral issues as well.

1.8.4 Policy Findings

Finding 9: Enduring unilateral dominance in cyberspace is neither realistic nor achievable by the United States.

In the event that conflict does occur, U.S. military doctrine seeks dominance in the relevant domains of conflict—that is, U.S. freedom of action in any domain of conflict (including cyberconflict) coupled with denying U.S. adversaries the same freedom of action.¹⁴ Dominance requires superiority in both offensive and defensive capabilities.

- Many cyberattack technologies are inexpensive and easily available to non-state actors, including individuals, and these technologies include some that are as capable of doing great harm as those available to governments. Much of the expertise needed to wield cyberattack weapons effectively is widespread. These points, discussed further in Chapter 2, suggest that the United States cannot maintain overall dominance in cyberattack capabilities for any extended period of time.
- With respect to cyberdefense, current trends in information technology development and deployment suggest that exploitable vulnerabilities will continue to be present in both civilian and military computer systems and networks of the United States. Thus, the U.S. information technology infrastructure is likely to remain vulnerable to cyberattack for the foreseeable future.¹⁵

Thus, cyberconflict is quite unlike the land, air, and maritime domains in which U.S. armed forces operate, and enduring unilateral dominance

¹⁴ According to the Joint Chiefs of Staff, joint force commanders are called upon to “seek superiority early in air, land, maritime, and space domains and the information environment to prepare the operational area and information environment and to accomplish the mission as rapidly as possible.” Joint Publication 3-0, *Joint Operations*, February 13, 2008, available at http://www.dtic.mil/doctrine/jel/new_pubs/jp3_0.pdf.

¹⁵ In addition, another nation may impose by decree cybersecurity measures on all information technology used by that nation or it may impose and enforce a strong separation between the information and information technology infrastructures for military and civilian use. Such a nation would likely have advantages in a cyberconflict with the United States, which does not do either of these things.

with respect to cyberconflict is not realistically achievable by the United States. This does not mean that the United States should refrain from developing cyberattack capabilities—only that it should not expect enduring advantage from such development.

Finding 10: The United States has much to lose from unrestrained cyberattack capabilities that are proliferated worldwide.

The United States is highly dependent on the capabilities afforded by ubiquitous information technology in every sector, both military and civilian. Consequently, the United States has much to lose from unrestrained cyberattack capabilities that are proliferated worldwide. (Some analysts also make the further argument that the United States would have the *most* to lose compared to any other nation—an assessment that is plausible but that depends on relative judgments about the dependence on information technology of various nations. The committee would not dispute that conclusion if it were accompanied by a defensible analysis, but it was not willing to make that assessment itself.)

In addition, comparing the as-yet-unproven utility of U.S. cyberattack against its adversaries to the demonstrated growing dependence of the United States on information technology, it is generally more important for the United States to be able to use information technology freely in pursuit of its national interests than for it to be able to deny adversaries the use of their own systems and networks. However, this conclusion does not rule out the possibility that cyberattacks by the United States will be an appropriate and useful action under some circumstances, although it does emphasize the importance of protecting the U.S. information technology infrastructure.

Finding 11: Deterrence of cyberattacks by the threat of in-kind response has limited applicability.

In general, deterrence of adversaries is the cornerstone of U.S. military strategy. Deterrence seeks to promote stability by persuading an adversary to refrain from taking aggressive actions against U.S. interests. Deterrence is based on two elements—punishment and denial. Deterrence by punishment threatens to inflict unacceptable costs on an adversary that takes aggressive actions. If he knows he will suffer such costs should he take such actions, he will refrain from taking them. Deterrence by denial seeks to deny the adversary success from his aggressive actions. If he knows his aggressive actions will not result in success, he will refrain from taking them.

As applied to cyberconflict, deterrence is complex. For the most part, defensive capabilities contribute to deterrence by denial, and attack

capabilities contribute to deterrence by punishment. Actions taken to strengthen important U.S. computer systems and networks promote deterrence by denial, but for a host of reasons described in Chapter 2 and in other reports,¹⁶ the gap between defensive capabilities and the adversarial cyberattack threat is large and growing today.

Deterrence by punishment is more likely to be an effective strategy against nations that are highly dependent on information technology, because such nations have a much larger number of potential targets that can be attacked. Nevertheless, even nations with a less technologically sophisticated national infrastructure are probably vulnerable to cyber-attack in selected niches.

A cyber aggressor also knows the time of his cyberattack, and can take action to mitigate the punishment that will follow his attack. The aggressor can take steps to invalidate the intelligence information on cyber targets that the defender has already collected on him, and thus can force the defender into either a non-selective retaliation or a retaliation delayed until new intelligence information can be collected. In the first case, the defender may not be willing to risk the large-scale escalation that might accompany a non-selective retaliatory cyberattack, and in the second case, the aggressor may have already achieved its objectives by the time a new retaliatory strike can be planned.

Perhaps most importantly, deterrence by punishment requires knowledge of an adversary's identity—anonymous adversaries cannot be punished. As noted in Chapter 2, today's information technology makes it easy for evildoers to act anonymously—and even in the event that new information technologies are developed with stronger authentication capabilities, there is always the risk that an authenticated computer could be improperly compromised to conduct aggressive action. On the other hand, an actionable degree of attribution might be possible by making use of non-technical information. Policy makers seeking absolute and unambiguous technical proof that a specific party is responsible for a cyberattack will almost certainly be disappointed in any real-life incident, and may ultimately be forced to rely on non-technical information more than they would prefer. 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.

Assuming that the adversary's identity can be known, there is no reason that a retaliatory cyberattack would necessarily be favored over a retaliatory kinetic attack. A variety of considerations might apply to choosing the retaliatory mode. For example, a "tit-for-tat" retaliatory

¹⁶ See, for example, National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

response against an adversary might call for a cyberattack of comparable scale against a comparable target. However, a threat to do so might not be credible if the United States has a great deal to lose from such an action, thus throwing doubt on the viability of an “in-kind” deterrence strategy. On the other hand, a near-peer competitor might well be deterred from launching a large-scale cyberattack by the knowledge that it too would have much to lose if the United States launched an in-kind counterattack.

If an access path is available to the adversary, it may be reasonable to use attack capabilities to neutralize an incoming cyberattack even if the identity of the adversary is not known. By developing capabilities to deny the adversary a successful cyberattack, the United States might be able to deter adversaries from launching at least certain kinds of cyberattack against the United States. Yet neutralization is likely to be difficult—destroying or degrading the source of a cyberattack may simply lead the adversary to launch the attack from a different source. Deterrence also relies on the adversary’s belief that the United States is indeed capable of neutralizing its attack—and such capabilities may well have to be demonstrated in order to induce that belief. But a demonstration may provide an adversary with ways of defending against those capabilities, and so the fragility of cyberweapons, noted in Chapter 2, may itself provide disincentives for the United States to provide such demonstrations. These disincentives may *raise* the thresholds at which the United States is willing to use those particular weapons. Thus, neutralization may be an appropriate response strategy, but whether a threat to neutralize an adversary’s attack is a reasonable basis for a strategy of deterrence through denial remains to be seen.

As for the tailored deterrence discussed in Chapter 9, that concept is premised on an understanding and a knowledge of specific adversaries. Indeed, it presumes that such knowledge is available *in advance* as the basis for tailoring a deterrence strategy against that particular adversary. But by definition, deterrence cannot be tailored to an adversary about whom nothing is known.

Against non-state parties, deterrence by punishment may be particularly ineffective, as noted in Section 9.3. First, a non-state group may be particularly difficult to identify. Second, it is likely to have few if any information technology assets that can be targeted. Third, some groups (such as organized hacker groups) regard counterattacks as a challenge to be welcomed rather than something to be feared. Fourth, a non-state group such as a terrorist or insurgent group might seek to provoke cyber retaliation in order to galvanize public support for it or to antagonize the public against the United States.

Finding 12: Options for responding to cyberattacks on the United States span a broad range and include a mix of dynamic changes in defensive postures, law enforcement actions, diplomacy, cyberattacks, and kinetic attacks.

Today, important information systems in the United States are subject to innumerable hostile actions on a daily basis from a variety of actors ranging from teenagers acting on their own to major nation-states.¹⁷ An important question for policy makers to address is thus, How should the United States respond to such attacks? And if or when the nature of cyberattacks changes in the future, how should it respond to those attacks?

Such questions cannot be addressed in the absence of specific facts. But it is important to understand that the United States has a multitude of options for responding to any given cyberattack, depending on its scope and character; these options include a mix of dynamic changes in defensive postures, law enforcement actions, diplomacy, cyberattacks, and kinetic attacks. Put differently, the United States is in no way obligated to employ an in-kind response to a cyberattack, even if an in-kind response may superficially seem most obvious or natural.

Some of the potential responses are less escalatory (e.g., changes in defensive postures), others more so (e.g., retaliatory cyberattacks or kinetic attacks). Implementing less escalatory responses would seem to require lower levels of authority than would more escalatory responses, and thus would be more easily undertaken.

1.8.5 Technical and Operational Findings

Cyberattack technologies are a relatively new addition to the technologies of warfare.

Finding 13: For many kinds of information technology infrastructure targets, the ease of cyberattack is increasing rather than decreasing.

Many recent reports have noted that the increasing use of information technology in existing and new infrastructure in the United States is increasing the vulnerability of that infrastructure. For example, *Toward a Safer and More Secure Cyberspace* notes that an increasing dependence on information technology applications in all walks of life has resulted in

¹⁷ See, for example, Dennis Blair, Director of National Intelligence, *Annual Threat Assessment of the Intelligence Community for the Senate Select Committee on Intelligence*, February 12, 2009, available at <http://intelligence.senate.gov/090212/blair.pdf>.

vulnerabilities being created faster than they can be found and fixed.¹⁸ Because the culture of information technology development in the United States does not promote security, new technologies, new architectures, and new applications result in new opportunities for attack. Old technologies and legacy systems also exhibit significant vulnerabilities because retrofitted security is often much less effective than security that is “designed in” from the start. The times required for defenders to repair security holes are long compared to the times required for attackers to develop new attacks. Many individuals and institutions do not know how to defend themselves because it is hard to do, and this is especially true of end users and small organizations.

These comments are also likely to be true for many other parties as well—to the extent that other nations are becoming dependent on information technology, there is no reason to suppose that they do not suffer from the same kinds of vulnerabilities. This is not to say that cyberattack on certain specific targets will not be very difficult, or that all cyberattacks can be assured of success with high probability. But on average and as argued in many reports,¹⁹ the gap between the attacker’s capability to attack many vulnerable targets and the defender’s inability to defend all of them is growing rather than diminishing.

Finding 14: Although the actual cyberattack capabilities of the United States are highly classified, they are at least as powerful as those demonstrated by the most sophisticated cyberattacks perpetrated by cybercriminals and are likely more powerful.

The cyberattack capability of a major nation-state (such as the United States) is almost certainly greater than that of the individual hacker or even the most talented cybercriminals. Such greater capability arises primarily from the resources available to nation-states rather than from fundamental differences in the base technologies available. A nation-state can draw on the services of its intelligence services and the funds in its national treasury, has enormous influence with the private sector companies over which it has jurisdiction, and is more than willing to bribe or extort to compromise a trusted insider if that is a cost-effective route to its objectives. In addition, it is entirely possible that certain technical problems have solutions that are today classified and thus not available to the

¹⁸ National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington D.C., 2007.

¹⁹ See, for example, National Research Council, *Information Technology for Counterterrorism: Immediate Actions and Future Possibilities*, The National Academies Press, Washington, D.C., 2003; and National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

world at large. In the domain of cryptography, it is known that the British Government Communications Headquarters (GCHQ; the UK equivalent of the National Security Agency) knew of public key encryption and in particular of the RSA algorithm for public key encryption several years before they were announced in the open literature.²⁰ Thus, one might reasonably presume that there may well be technical approaches to various forms of cyberattack that are known, at present, only on the “inside.”

The open literature documents a variety of sophisticated cyberattacks and cyberexploitations that have been used by criminals, and tools for these activities available to them. It is thus reasonable to posit that some of the tools available to nation-states are more sophisticated versions of criminal tools, that the associated procedures and practices are also more sophisticated versions of social engineering, and that the intelligence services of nation-states have greater capabilities with respect to cyberattacks that depend on some kind of close access. Put differently, the cyberattack capabilities of a major nation-state are at least as capable as those of sophisticated cybercriminals from a technical standpoint, and the attacks undertaken by such parties have been sophisticated indeed.

The comments above notwithstanding, non-state actors have certain advantages over nation-states. Non-state actors are known for their ability to act and react more nimbly. Neither terrorists nor criminals are subject to the often-ponderous processes of governmental oversight, which suggests that they may be able to move faster to take advantage of emergent opportunities for cyberattack (e.g., the approvals needed to conduct a cyberattack are likely to be fewer than in the U.S. government). Nor are they likely to adhere to either the letter or spirit of the laws of armed conflict in conducting their cyberattacks, which suggests that their planning is likely to be simpler and face fewer constraints (e.g., they can avoid the need to minimize collateral damage). And in a search for technical expertise and talent, they can often offer financial compensation that far exceeds anything that the U.S. government can legitimately offer its employees or troops. Whether such advantages offset the nation-state’s superiority of resources and access regarding actual operations and the best use of available capabilities is not clear.

Finding 15: As is true for air, sea, land, and space operations, the defensive or offensive intent motivating cyber operations in any given instance may be difficult to infer.

²⁰ Peter Wayner, “British Document Outlines Early Encryption Discovery,” *New York Times*, December 24, 1997, available at <http://www.nytimes.com/library/cyber/week/122497encrypt.html#1>.

This report distinguishes between different kinds of cyber actions (that is, cyber actions with different effects) and different intents (whether an action is carried out with offensive or defensive intent). A useful analogy is military mines. From a technological standpoint, a mine is an explosive device designed to explode (for example) on contact with a vehicle or a vessel. But it can be used for both defensive and offensive purposes. U.S. mines in the Korean demilitarized zone are intended to slow a North Korean attack, and their deployment is thus defensive in intent. U.S. mines in Nicaraguan ports were intended to contribute to the economic isolation of Nicaragua, and thus their deployment was offensive in intent.

Similarly, a U.S. cyberattack against a specific Zendian computer system may be conducted in order to stop an attack on U.S. systems emanating from that Zendian system (a defensive use), or it may be conducted in order to cripple a military computer system in anticipation of a U.S. kinetic attack on Zendia (an offensive use). Such issues affect perception across national boundaries as well—what the United States regards as a defensive action another nation may regard as an offensive action. And both perceptions would have some factual basis.

Furthermore, and as noted in Section 9.2.2, it may be more difficult to discern or assess intent when cyberattack is involved than when traditional military forces are involved. From a policy perspective, this point regarding the difficulty of inferring intent is significant when the United States is the target of cyberattacks as well as when it conducts cyberattacks. The strategic significance and societal effect of a cyberattack on the United States originating with an overly curious teenaged hacker in San Diego or in Mexico City is not the same as one originating from Zendia's 342nd Information Operations regiment, although in the initial stages of a cyberattack on the United States, it may not be entirely clear which of these parties is behind it. At the same time, if an adversary is uncertain about the intent behind a cyberattack emanating from the United States, its reaction may well be hard to predict.

Finding 16: Certain cyberattacks undertaken by the United States are likely to have significant operational implications for the U.S. private sector.

The private sector owns and operates much of the infrastructure through which certain cyberattacks might be transmitted and also has a significant stake in the continuing operation of that infrastructure, in particular the Internet. Thus, cyberattacks launched through the Internet may well have implications for and impacts on other non-military national interests.

It is not new that military decision makers must consider the impact of such decisions on other parties; for example, military decision makers have long known that reducing the availability of GPS satellites could have a major impact on non-military transportation, and efforts to jam adversary radars might impact non-military communications. However, in many such instances, the impacts could be spatially localized, e.g., by reducing the availability of GPS satellites only in the theater of conflict. Furthermore, because the adversary has been assumed to be confined to a particular geographic theater, adversary reactions to U.S. offensive operations have been confined as well—thus allowing non-military national activities outside the theater to be pursued under more or less normal conditions.

However, the Internet portion of cyberspace is entirely shared between military and civilian uses and between the United States and adversaries. Thus, the U.S. private sector must be prepared to deal with the consequences should the United States take actions that provoke in-kind counterattack by an adversary. In addition, the United States must consider the possibility that a cyberattack of its own, carried over the Internet, might be detected by U.S. Internet service providers carrying that traffic—and then shut down in the (mistaken) belief that it was an attack being carried out by hackers or another nation.

Lastly, U.S. cyberattacks that are directed against globally shared infrastructure supporting the private sector might have deleterious “blow-back” effects on U.S. private sector entities. Such effects might be direct, in the sense that a U.S. cyberattack might propagate to harm a U.S. firm. Or they might affect the supply chain of a U.S. firm—a node in Zendia might support communications between a key U.S. firm and a supplier firm in Ruritania as well as military communications in Zendia, and a disabling cyberattack on that node might leave the U.S. firm without the ability to order goods from the Ruritanian firm.

Finding 17: If and when the United States decides to launch a cyberattack, significant coordination among allied nations and a wide range of public and private entities may be necessary, depending on the scope and nature of the cyberattack in question.

Significant amounts of coordination with multiple parties may be required if and when the U.S. government contemplates the use of cyberattack. Although cyberattacks that are narrowly focused on highly specific objectives may not have much potential for interfering with other ongoing cyber operations initiated by other parties, a sufficiently broad cyberattack might indeed interfere. In such cases, it may be necessary to

coordinate among a number of parties, including various U.S. government agencies and allied nations. All of these parties may have various cyber operations underway that might interfere with a U.S. cyberattack on an adversary. In addition, these agencies and nations would likely benefit from the strengthening of their defensive postures that could occur with advance notice of a possible in-kind response. The same considerations apply to private sector operators of information infrastructure that would be likely targets of an adversary's in-kind response to a U.S. cyberattack and for which advance notice of cyberattack would be helpful in strengthening their defensive posture, although selective notification for operators of U.S. information technology infrastructure may raise issues of discrimination comparable to those that led to the State Department's adoption (after the Pan Am 103 bombing) of a policy of not warning government employees of a terrorist threat without making a general public announcement of the threat. Certain kinds of cyberattack may require the cooperation of various vendors—e.g., virus attacks that depend on disabling antivirus protection supplied by U.S. or foreign vendors or denial-of-service attacks requiring increased bandwidth supplied by U.S. or foreign Internet service providers.

Finally, the United States is likely to have in its midst “patriotic hackers”—U.S. citizens and others who are strongly motivated to take direct action in putative support of an overt U.S. confrontation with another nation. These individuals—private citizens with some skills in the use of cyberattack weapons—might well launch cyberattacks on an adversary nation on their own initiative, that is, without the blessing and not under the direction or control of the U.S. government. Such actions might interfere tactically with operations planned by the U.S. government, and strategically they might be misinterpreted by the party being attacked as intentional U.S. actions and thus complicate the conduct of diplomatic action.

Thus, the U.S. government would have to be prepared to discourage their actions using all legal means at U.S. disposal (e.g., through law enforcement authorities seeking to enforce the Computer Fraud and Abuse Act against these patriotic hackers) and would have to anticipate in its planning the actions that were not discouraged. Such means are not limited to prosecution (which would almost surely require a time scale much longer than that of a U.S. cyberattack); other legal means are often available to shut down the operational capability of a patriotic hacker, including arrest, seizure of the computer involved, disconnection from the Internet service provider that the hacker uses, and so on.

In extreme cases, the agency conducting the cyberattack might also find it necessary to conduct a cyberattack to neutralize the civilian system involved in this unhelpful hacker activity. Clear standards, thresh-

olds, and approval requirements would be necessary if such action were contemplated, and higher authority would have to consider a variety of questions. What sort of “interference” with a U.S. cyberattack is enough to justify an attack on a civilian U.S. system? What sort of circumstances would warrant such an attack? Need legal methods be attempted first? What level of certainty must exist about the involvement of the site before it can be attacked? Who must give the approval?

Finding 18: The outcomes of many kinds of cyberattack are likely to be more uncertain than outcomes for other kinds of attack.

Although planners for any kind of attack, kinetic or cyber, must take into account many uncertainties about the characteristics of the target and the environment around it, the intelligence information needed for a successful cyberattack (e.g., details of cabling between two systems) is often difficult to obtain through traditional methods such as remote photo reconnaissance. Such uncertainties can increase significantly the likelihood of unintended and/or unanticipated consequences. By contrast, many of the uncertainties in kinetic targeting can be calculated and bounded, and most of the remaining uncertainties relate to matters such as target selection and collocation of other entities with the intended target.

These comments should not be taken to imply that the mere presence of uncertainty renders cyberweapons inherently unusable. In some cases, operational or policy goals may require “taking a chance” even if the uncertainty of a given cyberattack’s effects is large. In other cases, the uncertainty inherent in a given cyberattack may not be significant from an operational or policy perspective. Moreover, the uncertainty associated with any and all cyberattacks is not necessarily large. A cyberattack might be designed to affect only a specific computer with a specific known serial number—such an attack would have few ill effects on any other computer system. A close-access cyberattack on a computer without electronic connections with the outside world is very unlikely to have effects in the outside world, as long as it remains isolated. A cyberattack using software agents exploiting vulnerabilities in Linux cannot necessarily exploit similar vulnerabilities on computers running Windows or Macintosh-OS systems. But greater intelligence efforts to resolve uncertainties are likely to be necessary to achieve levels of confidence equivalent to those that generally characterize kinetic attacks—and such efforts may in some cases take long enough to render the use of cyberattack moot.

Finding 19: Early use of cyberattack may be easy to contemplate in a pre-conflict situation, and so a greater degree of operational

oversight for cyberattack may be needed compared to that for the use of other options.

It is easy to see how policy makers might regard cyberattack as a desirable option when coercive measures are needed. Cyberattack can be portrayed as an instrument that is easy, simple, temporary, reversible, non-lethal, and non-risky for the United States to use. But although it is possible to imagine that such attributes might characterize some cyberattacks, the committee believes that such claims should generally engender a certain degree of skepticism among policy makers.

For example, a cyberattack might be regarded as a minor step. Although it is not new that “small” activities in a preconflict situation may have large consequences,²¹ the operational footprint left by cyberattack activities is small, a fact that tends to render activities related to this area less visible to senior decision makers. Given the fact that cyberattack may have strategic significance (perhaps inadvertently),²² senior military commanders (for example) will need to take special care to maintain situational awareness and affirmative control of their own forces under these circumstances and will need to exercise a greater degree of oversight than might be necessary if only conventional military forces are involved. (Of course, they also need to maintain awareness of adversary forces.)

Similar considerations apply to those responsible for making decisions about covert action. From a technical standpoint, cyberattack is an instrument that is well suited to covert action because of the inherent deniability of a cyberattack and the ability to conduct such an attack without “boots on the ground” (and thus without placing U.S. or other friendly lives at risk). This point is *not* intended to comment on the desirability of covert action as an option for U.S. decision makers—only that should covert action be determined to be desirable and in the national interest, policy makers are likely to be drawn to cyberattack as a preferred methodology for implementing such action. Accordingly, all of those responsible for exercising oversight over covert actions up the entire

²¹ For example, during the Cuban Missile Crisis, a U-2 reconnaissance aircraft on a “routine air sampling mission” over Alaska went off course and flew into Soviet airspace. The Soviet Union scrambled fighters to intercept the airplane, and the United States scrambled fighters to provide cover for the U-2. These U.S. fighters had been armed with nuclear air-to-air missiles. Upon hearing this news, Secretary of Defense Robert McNamara expressed grave concerns that the U-2 flight could have been interpreted as the prelude to a U.S. nuclear strike on the Soviet Union. See Max Frankel, *High Noon in the Cold War: Kennedy, Khrushchev, and the Cuban Missile Crisis*, Random House, New York, 2005.

²² Consider the possibility that a nuclear-armed nation might respond with the use of nuclear weapons to a major cyberattack (i.e., one with major societal consequences), as discussed in Section 10.3.

chain of command must be cognizant and aware of the risks, benefits, and uncertainties that they entail, whether they involve the use of cyberattack or other instruments. The possibility of using cyberattack as a means of covert action may also tempt decision makers to think that they might conduct a covert action with very little chance of detection—and therefore might lead to an inclination to intervene simply because the risks of detection are seen as lower.

Finding 20: Developing appropriate rules of engagement for the use of cyberweapons is very difficult.

Rules of engagement (ROEs) specify for military personnel the circumstances under which they can use their weapons and the authority required for doing so. Most importantly, ROEs are supposed to be developed *prior to* the need for use of their weapons, so that operators have proper guidance under operational circumstances. This fact means that various contingencies must be anticipated in advance, and of course it is difficult to imagine all possible contingencies before any of them happen.

Although ROEs normally are not specific to individual weapons systems, the presence of weapons or tools for cyberattack may be problematic. When cyberattack may be used, ROEs must be developed to cope with the fact that several dimensions of cyberattack span a wide range. A cyberattack may be non-lethal, or it may be destructive on a society-wide scale. The impact of a cyberattack can be easily predicted in some cases and highly uncertain in other cases. The set of potential targets that may be adversely affected by a cyberattack is quite large, and likely larger than the corresponding set of potential targets for other weapons. A cyberattack conducted for offensive purposes may well require authorization from higher levels of command than would a technically similar cyberattack conducted for defensive purposes. The adversary might not react at all to a cyberattack, or it might react with nuclear weapons. The adversary might be a solo hacker or a well-funded nation-state. It is thus unrealistic to try to craft a single ROE that attempts to cover all uses of cyberattack. Rather, it will be necessary to tailor an array of ROEs that are applicable to specific kinds of cyberattack and for likely specific circumstances. And it will be at least as difficult to craft ROEs for missions involving cyberattack as for missions involving other kinds of weapons.

As an illustration of the complexity of developing ROEs in a specific situation involving cyberattack, consider some of the issues in developing, in advance, military ROEs for active threat neutralization—under what circumstances governed by what authority might a counter-cyberattack be launched to neutralize an immediate or ongoing threat?

- *Who should have influence on the development of ROEs for active threat neutralization?* It is obvious that the agency conducting a counterattack should have input (likely the DOD or the intelligence community). But other agencies (notably the Departments of Homeland Security, State, Justice, and Commerce) may have equities at stake as well. And although it makes little sense for Congress to be involved in approving rules of engagement in detail, Congress should have mechanisms for being kept informed of the general circumstances under which the U.S. government does undertake active threat neutralization.

- *How, if at all, are the intent and the identity of a cyberattacker relevant?* If the cyberattacker is determined to be a nation-state, does it increase or decrease the appropriateness of a neutralization effort? (If it depends on other factors, what other factors?) And if intent is relevant, how is the intent of the cyberattacker to be ascertained? Suppose that the proximate nodes involved in an attack can be identified but they are likely innocent parties who have been compromised—is it appropriate to neutralize the threat emanating from their systems?

- *How does the proportionality principle apply to active threat neutralization?* Proportionality requires that the value of neutralizing the threat be outweighed by the likely collateral damage of the counterattack. How is the likely collateral damage to be estimated, especially if the response is automated and launched without human analysis or intervention? Or does a proper proportionality analysis *require* human intervention before such a response is launched?

- *How far down the chain of command should delegation of authority to launch an active threat neutralization be carried?* For example, although the commander of the U.S. Strategic Command has the authority under standing rules of engagement to conduct a response action, it is unlikely (though possible) that he must himself approve the action. It is more likely that the authority to do so is further delegated to other parties down the chain of command. But since a response action is a serious thing, there must be limits (not known to the committee) to how far this authority is delegated. (An automated response, as proposed by the U.S. Air Force's Concept of Operations for its Cyber Control System, would represent the ultimate in delegation of authority.)

- *What level of impact (among other factors) must an incoming cyberattack threat achieve in order to justify an active threat neutralization?* The standard used by the U.S. Strategic Command is that an incoming cyberattack must have a material impact on the DOD's ability to perform a mission or to carry out an operation, and that cyberattacks that merely cause inconvenience or that are directed only at intelligence gathering do not rise to the threshold of warranting such a response. For example, a cyberattack on the command and control system for Navy ballistic missile submarines

might warrant an active threat neutralization, but a cyberattack on the administrative computers of the U.S. Strategic Command might not.

- *How should the scope, duration, and intensity of a neutralization action be calibrated?* The intent of neutralization is to stop an incoming attack. But the scope, duration, and intensity of a response may relate to one's confidence in actually achieving an effective neutralization, as well as to the collateral damage that may be incurred. In addition, political reality may dictate that only a commensurate response (i.e., a response that inflicts a similar amount of harm on an adversary) is possible—how might this requirement square with effectiveness in stopping the attack?

A further level of complication in developing rules of engagement is that the factors above cannot be assessed independently. For example, the authority needed to launch an active threat neutralization may depend on the identity of the attacker—perhaps local authority would be needed if the attacker were a teenager in Zendia, but perhaps the personal authority of the commander of U.S. Strategic Command would be needed if the attacker were the 418th Zendian Information Operations Brigade. Perhaps higher-level authority would be needed if more collateral damage were possible.

The difficulties in formulating appropriate rules of engagement, and of different human beings interpreting these rules in a manner consistent with the intent in formulating them, suggest that there may well be differences between what is intended and what is actually done—and furthermore that these differences reflect an enduring reality of the way such processes operate.

1.8.6 Organizational Findings

Finding 21: Both the decision-making apparatus for cyberattack and the oversight mechanisms for that apparatus are inadequate today.

Adequate policy decision making and oversight require a sufficient base of technical knowledge relevant to the activities in question, an organizational structure that enables decision making and oversight to take place, and information about activities that are actually undertaken under the rubric of policy.

Cyberattack is a relatively new addition to the menu of options that policy makers may exercise, and there are few precedents and little history to guide them today. The infrastructure and resources needed to conduct such activities, and the activities themselves, are by their nature less visible than those associated with more traditional military, intel-

ligence, or law enforcement activities. They do not fit into standard categories—the weapons involved initially act in a non-lethal manner, even though they may have subsequent effects that are lethal or destructive; the activities for which they are suited go far beyond just surveillance or just covert action; and they are shrouded in secrecy. In many cases, budgets to acquire cyberattack capabilities are likely small compared to the budgets for major weapons acquisition programs. The technical knowledge needed to conduct informed oversight is not widespread, and the importance of cyberattack as a possible option for policy makers is not widely appreciated. Procedures for informing potentially relevant policy makers in both the executive and the legislative branches appear to be minimal or non-existent.

To illustrate the committee's concerns, consider the delegation of authority to the commander of the U.S. Strategic Command for conducting an active threat neutralization (a limited and specific form of active defense) to protect military computer systems and networks whose mission performance has been compromised by a cyberattack. The implications of such an action conducted against computer systems or networks outside U.S. borders may range beyond strictly military ones, especially if the potential for unintended consequences is taken into account. This is not to say that all active responses have such potential, or that any active response will necessarily have unintended consequences. But absent mechanisms for factoring in diplomatic or political considerations, the committee is concerned about a decision to conduct an active threat neutralization that takes into account only military or local tactical considerations of protecting the mission capability of U.S. military networks.

With such factors in play, an adequate organizational structure for making decisions and exercising oversight has not emerged, and much of the information relevant to conducting oversight is unavailable. As a result, government and society at large are neither organized nor prepared to think about the implications of cyberattack as an instrument of national policy, let alone to make informed decisions about them.

Finding 22: The U.S. Congress has a substantial role to play in authorizing the use of military force, but the contours of that authority and the circumstances under which authorization is necessary are at least as uncertain for cyberattack as for the use of other weapons.

One important missing element—conspicuous in its absence—in the decision-making apparatus of the U.S. government is the role that the Congress does or should play in decisions related to cyberattack. As noted in Chapter 6, Congress has an important authorization role regarding

the use of military force under many circumstances, although the limits of that authority are the subject of much dispute between the executive and legislative branches. If the necessity of congressional authorization for the use of traditional U.S. military forces is disputed as it has been in recent U.S. history, consider the conundrums that could accompany the use of weapons that are for all practical purposes covert and whose “deployments” would be entirely invisible to the public or even to most uniformed military personnel.

In general terms, the use of cyberattack raises the same sorts of issues as other instruments of warfare such as frigates and cruise missiles. When does the President have inherent authority to act regardless of what Congress says or does? When must the President obtain congressional approval before acting? When can Congress define the standards and procedures that limit what would otherwise be plenary presidential authority? Nevertheless, cyberweapons raise particularly difficult issues in this context (as do certain kinds of non-cyberweapons), because of the need for speed in using such weapons (e.g., because of a target’s transience), the risk of unintended and unknown consequences, and the lack of visibility of their use.

The committee refrains from making a finding on the boundaries between presidential and congressional authorities in this area, but notes the existence of certain limiting cases on both sides of this debate. In one limiting case, the committee believes it would be broadly accepted that presidential views of executive branch powers notwithstanding, congressional authorization is required for the United States to launch a large-scale cyberattack against another nation with the intent of shutting down the essential civil services of that nation—transportation, electric power, financial services, and so on—if the attack were contemplated as a first use of coercive or aggressive action against that nation.

In another limiting case on the other side, the committee believes that there are certainly some circumstances under which some kind of cyberattack might be launched without explicit congressional authorization, just as certain kinds of military force can be used under some circumstances without such authorization. The canonical example of the latter is the use of force in self-defense—if U.S. military units are attacked, standing rules of engagement generally permit the use of lethal force against the attacking party.

However, in the vast area of possible circumstances in between these two limiting cases in which the United States might contemplate a cyberattack, the lines are most unclear, and the committee is explicitly silent on those lines.

A variety of factors may influence whether a given situation falls above the line requiring congressional authorization or below the line.

Possibly relevant factors include the scale of the cyberattack contemplated, the target of the cyberattack, and the circumstances that define “first use.” One particularly problematic issue is the possibility of escalation and unanticipated effects, in which cyberattacks that do not require congressional authorization might evolve into cyberattacks that do. (Unanticipated effects are, by definition, unintentional, although they might well not be perceived by the attacked party as unintentional.) The escalation issue is also present in a non-cyber context, and is indeed what the War Powers Resolution was intended to prevent, but as discussed in Section 6.2.1, the cyber dimension of the issue significantly increases the complexity of the problem.

Finally, the committee calls special attention to the fact that congressional concerns about asserting authority over the use of military forces are generally at their maximum when U.S. military forces are placed directly in harm’s way—that is, when U.S. casualties may be the result of direct combat. Cyberattacks launched by the United States are highly unlikely to place U.S. forces at direct risk, and indeed would in general be easy to undertake with minimal public visibility. Thus, explicit mechanisms to provide relevant information to the appropriate congressional parties are essential if Congress is to know if and when it should be involved.

1.9 RECOMMENDATIONS

U.S. acquisition and use of cyberattack capabilities raise many issues in need of broad understanding and deserving of extensive and widespread national conversation and debate. One set of committee recommendations focuses on fostering that debate. A second set of recommendations focuses on operational needs.

A caution to the reader: For the most part, the recommendations below are formulated as advising that “the U.S. government should do X or Y.” This formulation violates a basic canon of making recommendations to policy makers, namely that the party viewed by the committee as responsible for taking action on a recommendation should always be made as specific as possible. However, consistent with Finding 21, the committee could not identify an appropriate entity within the U.S. government to take action, and indeed as this report is being written, the U.S. government is trying to decide how best to organize itself internally to deal with the implications of cyberattack as an instrument of national policy.

1.9.1 Fostering a National Debate on Cyberattack

Recommendation 1: The United States should establish a public national policy regarding cyberattack for all sectors of government, including but not necessarily limited to the Departments of Defense, State, Homeland Security, Treasury, and Commerce; the intelligence community; and law enforcement. The senior leadership of these organizations should be involved in formulating this national policy.

As noted in Chapter 6, the DOD Information Operations Roadmap of 2003 recommended that the U.S. government should have a declaratory policy on the use of cyberspace for offensive cyber operations. As the committee has been unable to find any such statement of declaratory policy, it concurs with and reiterates this call. At a minimum, such a policy would involve the DOD, the intelligence community, and law enforcement agencies, and would address the following questions:

- For what purposes does the United States maintain a capability for cyberattack?
- Do cyberattack capabilities exist to fight wars and to engage in covert intelligence or military activity if necessary, or do they exist primarily to deter others (nation-states, terrorist groups) from launching cyberattacks on the United States?
- If they exist to fight wars, are they to be used in a limited fashion? Under what circumstances would what kinds of cyberattack be launched?
- What legal regimes are relevant to different levels of cyberconflict?
- How and when is cyberconflict to be stopped?
- To the extent that cyberattack is part of the U.S. deterrent posture, how can its use be established as a credible threat?
- What, if any, role do cyberattack capabilities have in law enforcement efforts directed against transnational criminal groups?

A clear statement of policy in this area would enable various government actors, and the private sector as well, to understand the constraints and limitations on using cyberattack for various purposes and to establish appropriate standards of behavior in this domain. Appropriate policy would provide important guidance for U.S. armed forces, intelligence agencies, and others in a domain in which international and national law may be inadequate to manage the full ramifications of using cyberattack.

For example, the United States could declare its commitment to abiding by the laws of armed conflict with respect to cyberattack. Such a posture could well affect the willingness of other nations to make similar declarations. Another related example concerns the national military strategy of the United States. As noted in Section 6.1.1, the *National Military Strategy of the United States*, published in 2004, indicated that the United States could respond using nuclear weapons to certain kinds of large-scale cyberattacks. Does this presumably authoritative statement of 2004 continue to reflect U.S. policy? If not, how does current policy differ? If so, is this an appropriate policy?

The new administration could undertake a review of cyberattack policy comparable to the nuclear policy review that new administrations often perform. Congressional hearings on this topic would also be useful in shedding light on government thinking about this topic.

The promulgation of a comprehensive declaratory policy would be a good first step for the government in becoming more forthcoming about its own thinking in this area and providing a benchmark for public discussion. But although the committee endorses the 2003 DOD recommendation regarding establishment of declaratory policy with respect to military equities, the committee goes further still. As noted in Finding 1, the committee believes that U.S. acquisition and use of cyberattack raises many important policy issues that go far beyond the Department of Defense. Such issues deserve an extensive and widespread national conversation and debate about how cyberattack might affect a broad spectrum of national interests.

The Departments of State, Homeland Security, and Treasury, and law enforcement agencies are thus included in Recommendation 1, even though they are not traditionally regarded as agencies with interests in cyberattack. The State Department is included because cyberattack has many international dimensions. The DHS and law enforcement agencies are included because tracing the ultimate source of an incoming cyberattack often requires the investigator to penetrate intermediate nodes, capture the relevant traffic, and then analyze it to determine the next node in the chain. Law enforcement authorities are also responsible for aspects of preventing or prosecuting cybercrime. The Department of the Treasury has responsibility for enforcing sanctions, and cyberattack may be relevant to the performance of this mission. In addition, implementation of Recommendation 10 may call for the establishment of an agency or a body with certain law-enforcement-like responsibilities that would also find some utility in conducting certain kinds of cyberattack.

Recommendation 2: The U.S. government should conduct a broad, unclassified national debate and discussion about cyber-

attack policy, ensuring that all parties—particularly Congress, the professional military, and the intelligence agencies—are involved in discussions and are familiar with the issues.

As noted in the Preface, the topic of cyberattack is highly classified within the U.S. government. Some aspects of the topic are classified with good reason—these include the fact of U.S. interest in a specific cyberattack technology (rather than the nature of that technology itself); fragile and sensitive operational details that are not specific to the technologies themselves (e.g., the existence of a covert operative in a specific foreign country or a particular vulnerability); or capabilities and intentions of specific adversaries. But the details of these areas are not particularly relevant to answering questions about declaratory policy, and thus secrecy even about broad policy issues serves mostly to inhibit necessary discussion about them.

Although implementation of Recommendation 2 would benefit both the private and public sectors of the nation as a whole, two stakeholder groups have particular significance. Both the U.S. Congress and the professional military/intelligence agencies need at least a basic understanding of the policy issues and their relationship to the basic technologies involved, but the broad classification of virtually all issues related to cyberattack is a significant barrier to the discharge of their responsibilities.

Recommendation 3: The U.S. government should work to find common ground with other nations regarding cyberattack. Such common ground should include better mutual understanding regarding various national views of cyberattack, as well as measures to promote transparency and confidence building.

The committee believes that most other nations are no farther along in their understanding of the key issues than is the United States. It is therefore important for the United States to begin to find common ground on this topic with allies, neutrals, and potential adversaries. In this context, “common ground” is not a euphemism for treaties or arms control agreements regarding cyberattack. It is rather a term that denotes a common understanding of its significance for policy—and common ground is important for allies and adversaries alike if misunderstandings are to be avoided.

Consultations with allies of the United States (such as the NATO countries) are likely to be the easiest to undertake. Such consultations should take two tracks—between the governmental entities that would be responsible for executing cyberattacks in these nations and between the cognizant policy decision makers. At the very least, those with opera-

tional responsibility for attack execution need to develop mechanisms for coordinating cyberattacks so that they do not interfere with each other. And policy makers must be able to discuss issues related to cyberattack in an informed manner, without having to learn about them in the middle of a cyber crisis.

As an example of such consultation, NATO established in March 2008 the Cyber Defence Management Authority, which will manage cyberdefense across all NATO's communication and information systems and could support individual allies in defending against cyberattacks upon request.²³ One press report indicates that "the Authority will also develop and propose standards and procedures for national and NATO cyberdefense organisations to prevent, detect, and deter attacks," but will focus on defense "whether an attack comes from state, criminal or other sources."²⁴ Similar efforts to reach common understandings regarding cyberattack (and on the relationship of cyberattack to cyberdefense) would be helpful as well.

Consultations with potential near-peer adversaries, or with the United Nations, are more politically fraught, especially with the Russian Federation seeking to delegitimize cyberattack entirely as a method of warfare. But Russian proposals on this topic are based on a Russian view of the topic, and it is worth understanding in some detail the sources of Russian concerns, even if the ultimate result is an agreement to disagree about basic premises and concepts. More generally, it would be helpful for all of the world's nations to understand the scope and nature of their interests where cyberattack is involved, and the only way to begin the process of developing understanding is to start consultations.

There are, of course, multiple forums in which to initiate consultations. Treaty negotiations are one possible forum, although U.S. policy makers may feel that such a forum grants too much legitimacy to an idea deemed by many in the U.S. government to be inconsistent with U.S. national interests. The UN Security Council itself could be another forum for discussions. NATO or G-7 ministerial discussions could be used to start consultations among allies.

The committee believes that greater mutual understanding and common ground should be sought on the following topics:

- *The scope and nature of cyberattacks, especially including those that would constitute a "use of force" and an "armed attack."* Given the overall lack

²³ NATO, "Defending Against Cyber Attacks: What Does This Mean in Practice?," March 31, 2008, available at http://www.nato.int/issues/cyber_defence/practice.html.

²⁴ See <http://www.computerweekly.com/Articles/2008/04/04/230143/nato-sets-up-cyber-defence-management-authority-in-brussels.htm/>.

of experience and history with cyberattack, such discussions would serve to provide common vocabularies and conceptual frameworks for addressing the issue in the future. What activities constitute a cyberattack? How might damage or harm from a cyberattack be assessed? What activities might constitute evidence of hostile intent? How should cyberexploitation and intelligence gathering be differentiated from cyberattacks? How, if at all, should exploitations for economic purposes be differentiated from exploitations for national security purposes? During discussions of these issues, no nation would have to acknowledge undertaking any of these activities since the intent would be the development of common conceptual frameworks.

- *Measures to promote transparency and to build confidence in the lack of aggressive intent.* By analogy to confidence-building measures in other domains (Chapter 10), the United States and other nations should seek to establish lines of communication between responsible and authoritative parties within their respective governments that would be able to account for or to deny suspicious cyber operations that might appear to be occurring at government direction. To make such communications meaningful, it would also be helpful for the nation involved to agree to cooperate in the investigation of such operations and/or to allow the victimized party to engage in self-help activities. Explicit agreement could be sought on what is required in order to “cooperate” on any investigation. For example, cooperation might require the nation hosting a network node involved in a cyberattack to provide forensic analysis of information on that node.

- *Building of informal relationships among key participants in both the public and the private sector.* One of the primary lessons of the Estonian incident of 2007 was the enormous value of relationships of trust among certain individuals in the Estonian government and top technical people from the various Internet service providers for Estonia and around the world. By exploiting these relationships, it was possible to take action to dampen the effect of the cyberattack against Estonia in a much shorter time than would have been possible if only formally sanctioned relationships between governments were available. Support and encouragement to develop such relationships should be provided by the governments involved.

- *Separation or identification of the computer systems and networks of military forces, the civilian population, and specifically protected entities (e.g., hospitals).* Much of the difficulty of adhering to the framework of the law of armed conflict in the context of cyberattack arises from the difficulty of distinguishing between valid military targets and other entities that are specially protected or are possible victims of collateral damage. It may be possible to develop mutually agreed methods or cooperative technical

means for cyberattackers to distinguish between these different categories to minimize inadvertent damage to non-military targets and ways to verify that these declared distinctions were being properly applied.

- *The significance of non-state parties that might launch cyberattacks, and how nations should respond to such attacks.* Today, the Convention on Cybercrime is the only international agreement on how nations should respond to cyberattacks, and is in essence an agreement to harmonize criminal law in this area and to facilitate law enforcement cooperation among the signatories. But as noted in the argumentation for Finding 7, the law enforcement framework operates in many cases on a time scale that is far too long to protect victims of cyberattack from harm.

1.9.2 Organizing the Decision-Making Apparatus of the U.S. Government for Cyberattack

Recommendation 4: The U.S. government should have a clear, transparent, and inclusive decision-making structure in place to decide how, when, and why a cyberattack will be conducted.

As noted earlier, the use of cyberattack in pre-conflict situations is likely to be tempting to policy makers. But because cyberattack can have far-reaching implications (at least in part because the actual scope of a cyberattack somehow gone awry may be much greater than that intended), senior policy makers should have a mechanism for ensuring that consultations take place with all stakeholders with equities that might be affected by a U.S. cyberattack in pre-conflict situations. At a minimum, it would appear that the Departments of Defense, State, and Homeland Security, and the law enforcement and intelligence communities would have to be involved in coming to terms with issues, such as advance coordination of a U.S. cyberattack that might lead to a cyberattack on the United States or to a determination that exploitation of adversary computers should (or should not) have priority over disabling or damaging them.

As an example of a question for which the U.S. government as a whole needs to establish an authoritative decision-making structure, consider cyberattack in the context of the dividing line between covert action and military activity. The U.S. Code defines covert action as “an activity or activities of the United States Government to influence political, economic, or military conditions abroad, where it is intended that the role of the United States Government will not be apparent or acknowledged publicly” (50 USC 413b(e)). At the same time, the U.S. code also defines any activity executed under control of the DOD chain of command as falling under the definition of a traditional military activity associated

with anticipated or ongoing hostilities, and such activity thus is not covert action subject to the findings and congressional reporting process.

The question of the boundaries between covert action and traditional military activities has been the subject of much discussion over the past several years (since the U.S. invasions of Afghanistan and Iraq). The findings and reporting process is often disliked by incumbent administrations, because it constrains an administration's ability to act freely and quickly and runs the risk of leaks that may reveal the existence of a covert action. On the other hand, many informed advocates of the process believe that the existence of such a process forces the executive branch to coordinate internal stakeholders and their equities, and also provides for necessary external review of actions that may be ill-advised from a broader public policy perspective.

The committee was not constituted to address this tension in its broadest formulation. But the stealthy operation and the difficulty of attribution associated with cyberattack weapons inherently makes them instruments of deniable action and may change the cost-benefit calculation in deciding whether a given covert action should be undertaken. Thus, the committee anticipates that this tension will increasingly manifest itself in a cyberattack context, and may push the boundaries of settled law and policy into uncharted territory. Accordingly, the committee believes that the issue is sufficiently important to warrant high-level attention from both the Administration and the Congress.

A second example of the need for an inclusive decision-making structure regarding cyberattack relates to active defense. As noted in Chapter 3, STRATCOM has the authority to neutralize a cyberthreat that compromises DOD mission effectiveness. Such authority is consistent with the traditional DOD standing rules of engagement that provide for force protection. Depending on the nature of the response action taken, however, a response may have strategic implications that go beyond force protection, even if the response action is limited in scope, effect, and duration. For example, if a cyberthreat is emanating from the military forces of a near-peer adversary, a response action may lead to escalation—especially if the response is not as controlled in execution as it was in planning or if the incident occurs during times of tension.

For such reasons, the committee believes that the decision to take such actions should be made at levels of authority high enough to weigh the various equities (military, diplomatic, and so on) appropriately. For example, the committee believes that the stakes of a neutralization cyberattack must be high enough (i.e., the damage being caused to computer systems and networks important and serious enough) and success likely enough to justify the political risks of launching a counterattack, such as the possibility that world opinion might not see U.S. cyberattacks under-

taken under the rubric of active defense as innocent acts of self-defense, even if they are. Such an assessment can be made only at the highest levels of government.

These points should not be taken to imply that the authority to conduct a neutralization response should not be delegated (though they do suggest that delegation should not go too far down the chain of command). Delegation with clear rules of engagement may be the only way to reconcile high-level decision making with the need for prompt response. Such rules would clearly establish the threshold at which a military mission is compromised and the constraints on the scope and nature of a neutralization response. For instance, one constraint might be conducting a neutralization response only when other methods²⁵ for responding to a cyberattack have proven (or will prove) ineffective. Another constraint might require that a neutralization response be limited in scope and as focused as possible on eliminating the threat in order to minimize the possibility of inadvertent escalation.²⁶ (Both of these constraints appear to be consistent with the rules of engagement described in Section 3.3 concerning possible DOD response actions for computer network defense.)

But because of the potential for erroneous response (Chapter 2 discusses the difficulties of attribution and attack assessment) and for inadvertent escalation (as described in Chapter 9), the committee is highly skeptical of the idea that delegation should include automated neutralization responses, a capability of interest to the U.S. Air Force (as noted in Box 3.5). Indeed, the authority for conducting a neutralization response should flow explicitly from higher authority, only after higher authority has considered all of the various equities in an integrated manner, and only after higher authority has reviewed and if necessary modified standing rules of engagement during times of crisis. Whether this description of the flow of authority in fact characterizes current rules of engagement for STRATCOM's authority to conduct response actions is not known to the committee.

A third example of the need for an inclusive decision-making struc-

²⁵ These other methods may include dropping connections, closing ports, asking Internet service providers to shut down nodes identified as being sources of the attack, diverting attack traffic to other locations, changing IP addresses, and so on.

²⁶ For example, consider two possible neutralization responses to a given botnet threat, wherein the botnet is controlled by a machine to which an access path has been established. One approach might be to launch a denial-of-service attack against the controller in order to prevent it from communicating with the bots it controls. Another approach might be to break into the controller to assume control of the botnet, and then issue orders to shut off the attack. Although the first method might be faster, it presumes that the attacked machine is dedicated to the controlling function, whereas in fact the machine in question might have other non-hostile functions whose termination might constitute an escalation.

ture can be seen in the fact that during active hostilities, a cyberattack conducted for tactical purposes might lead to opportunities whose exploitation would have strategic significance. For example, consider a cyberattack on the command and control network in the nationwide Zendian air defense system. In the process of exploring the network, corrupting data, and issuing confusing or damaging commands, U.S. operators might stumble onto a communications link with the Zendian national command authority (NCA). Exploitation of that link might enable the United States to penetrate the command and control network of the Zendian NCA—but a decision to do so should not be made by operators and commanders on the ground but rather by higher U.S. authorities. Thus, mechanisms must be established to provide such information up the chain of command when necessary, and other mechanisms established to act on such information should it be made available.

Recommendation 5: The U.S. government should provide a periodic accounting of cyberattacks undertaken by the U.S. armed forces, federal law enforcement agencies, intelligence agencies, and any other agencies with authorities to conduct such attacks in sufficient detail to provide decision makers with a more comprehensive understanding of these activities. Such an accounting should be made available both to senior decision makers in the executive branch and to the appropriate congressional leaders and committees.

Whether or not cyberattack falls into the category of covert action, it appears that even within the executive branch, knowledge of the actual cyberattack activities of the United States is highly fragmented. An authoritative source, updated periodically, that documents the extent and nature of such activities and provides analyses of their impact and/or significance would help senior decision makers within the executive branch and Congress in carrying out their authorization and oversight responsibilities.

The committee expects that such a compendium would be highly classified, as it would likely reveal many sensitive details regarding actual U.S. capabilities and actions. For understanding policy and for exercising oversight, such an accounting would describe the purposes served by any given cyberattack, the intended target(s), the outcome, the difficulties encountered in conducting the attack, the rules of engagement relevant to that cyberattack, and both the anticipated and the actual value of the attack in serving U.S. national interests. If necessary, exemptions to such reporting for extremely sensitive operations might be modeled on those

in the statute on covert action providing for more limited “Gang-of-Eight” reporting.²⁷

One approach to collecting the information would be for cyberattacks to be reported more or less contemporaneously to the National Security Council, which would compile and analyze the information and then distribute it when required to do so. This approach also has the advantage of informing senior executive branch decision makers of potentially significant events that might affect their activities and decisions in other domains (e.g., if undertaken in the middle of a crisis, an inappropriately timed cyberattack might have diplomatic repercussions).²⁸

Also, consistent with Finding 22, the committee recommends the establishment of mechanisms to promptly inform the appropriate parties in Congress before the United States launches significant U.S. cyberattacks against other powers or entities or promptly thereafter. “Promptly” should be understood to refer to a time scale shorter than or comparable to those required by the War Powers Resolution for introducing U.S. armed forces into hostilities.

Finally, the committee recognizes that many definitional issues remain to be worked out. It is the committee’s recommendation that a reportable cyberattack be defined as one that was initiated with the intent of altering, disrupting, deceiving, degrading, or destroying adversary computer systems or networks or the information and/or programs resident in or transiting these systems or networks immediately or in the future. For example, reasonable people might disagree over whether cyberexploitations should also be included, but the goal is for responsible senior decision makers to have a reasonably comprehensive view of the cyberattack-related activities of the U.S. government.

1.9.3 Supporting Cyberattack Capabilities and Policy

Recommendation 6: U.S. policy makers should judge the policy, legal, and ethical significance of launching a cyberattack largely on the basis of both its likely direct effects and its indirect effects.

²⁷ “Gang-of-Eight” reporting refers to the requirement to report only to the chair and ranking minority member of the House and Senate Select Committees on Intelligence, the Senate majority and minority leaders, and the Speaker of the House and the House Minority Leader. Reporting to the “Gang of Eight” meets the legal requirement for presidential briefing to Congress for certain selected intelligence activities.

²⁸ In this regard, executive branch notification might be regarded as being analogous to notifying the secretary of defense about all missile test launches. The intent of this long-standing rule was not that the secretary had to approve such launches but rather that the secretary should know if a launch was going to occur in the middle of other events or during a crisis.

As noted in Finding 5, the consequences of a cyberattack may be both direct and indirect—and both must be taken into account in determining appropriate courses of action. Cyberattacks cannot be assumed to be of lesser consequence simply because they are primarily non-kinetic attacks on computer systems or networks.

This point is especially relevant in considering responses to a crisis or an incident in which a forceful U.S. response is desired. Because a cyberattack may appear to be an action short of a “real” military deployment or response if only direct effects are considered, and in any event would be unlikely to place U.S. forces directly in harm’s way, policy makers may be unduly tempted to take such an action unless they consider the cyberattack’s indirect effects as well.

More generally, the difficult legal and ethical policy issues regarding the appropriateness of using cyberattack seem to arise mostly in a prekinetic situation, where traditional armed conflict has not yet arisen (and may never arise). In this context, decision makers must determine whether a cyberattack would be equivalent to “the use of force” or “an armed attack.” Effects-based analysis provides one criterion for such a determination—equivalence would be determined by comparing the scale of death and/or destruction that would result from a cyberattack (taking into account both direct and indirect effects) to that which would result from a use of kinetic force.

As for the situation in which a “kinetic” conflict has already broken out, cyberattack is just one more tactical military option to be evaluated along with other such options—that is, when U.S. military forces are engaged in traditional tactical armed conflict and except in extraordinary circumstances, there is no reason that any non-LOAC restrictions should be placed on the use of cyberattack vis-à-vis any other tactical military option. Thus, if a given tactical operation calls for attacking a certain target, LOAC questions about necessity, proportionality, and distinction must be asked about the use of cyberattack, the use of special operations troops, and the use of a cruise missile—and attacks that do not satisfy LOAC constraints may not be used. (Needless to say, both direct and indirect effects must be considered in this analysis, and uncertainties in the answers to these questions must be taken into account as well.)

The extraordinary circumstances mentioned above relate to instances in which U.S. military forces might be contemplating actions with strategic significance. For example, a cyberattack on an adversary satellite might have tactical benefits, but the use of a cyberattack for this purpose should be considered just as carefully as the use of a direct-ascent missile or a ground-based laser. The latter decision today would not be the sole province of the commander in the field, but would likely involve the National Command Authority directly, and so should the former. Com-

manders in the field should not be tempted by the seeming ease or low profile of cyberattack to use such an option when other options would not be used.

Finally, Recommendation 6 should not be taken to mean that only effects are relevant to a policy, legal, or ethical analysis of any given cyberattack. The committee recognizes, for example, that the intent with which a cyberattack is carried out may well be relevant to such analysis, though the attacker's intent may be largely irrelevant to its effects. Indeed, the DOD standing rules of engagement (mentioned in Section 3.3) obligate military commanders to "defend that commander's unit and other U.S. forces in the vicinity from a hostile act or *demonstration of hostile intent*." The party responsible for the attack is also a relevant factor—it matters whether the responsible party is a nation-state, terrorist group, criminal organization, hacker, or a careless graduate student. Thus, a cyberattack launched by a terrorist group affecting a small number of important national security computer systems may well be regarded as a more hostile act than a cyberattack launched by a careless graduate student affecting millions of systems around the world (including some national security computer systems)—and a national response should account for such differences.

Recommendation 7: U.S. policy makers should apply the moral and ethical principles underlying the law of armed conflict to cyberattack even in situations that fall short of actual armed conflict.

As noted in Chapter 7, the law of armed conflict—specifically *jus in bello*—does not pertain to the behavior of military forces in situations that fall short of actual armed conflict, and the relevant international law under such circumstances is poorly developed at best. Nevertheless, the committee believes that U.S. policy makers should apply the moral and ethical principles underlying the law of armed conflict *jus in bello* (proportionality, necessity, distinction, and so on) to cyberattack even if the use of cyberattack is contemplated for situations that fall short of actual armed conflict.

The application of these principles would be particularly relevant in two situations:

- *Covert actions involving cyberattack.* (As noted in Chapter 4, traditional U.S. interpretations of the laws of armed conflict require covert action, whether or not it involves violent activities, to be conducted consistent with LOAC's requirements.)
- *Periods of heightened tension,* during which combatant commanders

may undertake some cyberattack activities for shaping the operational environment to facilitate later employment of other activities (as noted in Chapter 3).

Recommendation 8: The United States should maintain and acquire effective cyberattack capabilities. Advances in capabilities should be continually factored into policy development, and a comprehensive budget accounting for research, development, testing, and evaluation relevant to cyberattack should be available to appropriate decision makers in the executive and legislative branches.

The committee believes that it would be unwise policy to eschew cyberattack under all circumstances. For those instances in which the use of cyberattack is warranted, the United States should have at its disposal the most effective and flexible cyberattack technologies and supporting infrastructure possible—systems that can operate on the time scales required, with the necessary command and control (including self-destruct when necessary and appropriate), guided by the best possible intelligence information, with a high probability of mission success and a low risk of collateral damage.

Accordingly, in addition to a robust and significant effort for research, development, testing, and evaluation to strengthen U.S. cyber defensive capabilities, the committee believes that the United States should continue to invest in the development and acquisition of effective and highly flexible cyberattack capabilities. In addition to providing operational utility, such capabilities may strengthen deterrence against cyber adversaries. Lastly, increased knowledge of cyberattack technologies will contribute to the knowledge base supporting development of improved defensive capabilities, assuming that mechanisms can be found to promote cross-fertilization among the researchers in the relevant areas.

If and when new policy emerges that calls for a deemphasis of cyberattack capabilities, the U.S. investment can be scaled back at that time. The committee recognizes precedents from history in which the momentum built up by a large-scale development and procurement plan made changes in policy more difficult to accomplish. Nevertheless, it believes that acquiring many kinds of cyberattack weaponry is relatively inexpensive compared to traditional large-scale weapons acquisition efforts, and thus policy changes would be easier to effect.

In addition, even if international agreements are made to restrict the use of cyberattack, nations must prepare for the possibility that non-signatories (e.g., non-state actors, or recalcitrant states) or “cheating” states will not abide by the provisions of any such agreement—and for the

United States to not be prepared to compete successfully in such a world is unacceptable.

Finally, it is important for the United States to have a comprehensive view of the effort among all of the relevant stakeholders to develop and acquire cyberattack capabilities. Some responsible party within the executive branch, perhaps an office within the Office of Management and Budget, should have a cross-agency view into overall amounts being spent on acquisition of cyberattack capabilities and the details of how individual agency budgets are being spent. Overall levels of spending and the relevant detail should be available, on a classified basis as necessary, to appropriate congressional decision makers. (Recommendation 8 is not a plea for centralized direction of the acquisition effort, but rather one for information to help policy makers understand the overall effort.)

Recommendation 9: The U.S. government should ensure that there are sufficient levels of personnel trained in all dimensions of cyberattack, and that the senior leaders of government have more than a nodding acquaintance with such issues.

The issues related to cyberconflict are quite complex. Conducting cyberattacks requires specialized expertise in operations, intelligence, and communications, as well as law and technology. Understanding policy related to cyberattack requires expertise in defense, intelligence, law enforcement, and homeland security, and in diplomacy, foreign relations, and international law. In short, the prospect of cyberconflict requires that considerable attention be given to professionalization of the involved workforce.

These needs contrast with the history of how today's thinking about cyberattack has evolved over the last few decades. The personal computers first introduced in the 1980s and then later the World Wide Web in the mid-1990s are the most visible signs of the information technology revolution that increasingly has affected all sectors of society, including the military. The possibility of information and information technology as the driver for a revolution in military affairs began to gain influence during this time, along with the notion of attacking an adversary's computers as an instrument of warfare. However, for the most part, that notion was confined to the grass roots of the military, and only recently has the thinking of senior military leadership begun to embrace such possibilities seriously.

Against this backdrop, the paucity of educational opportunities in this domain for senior leadership, the professional military, the diplomatic corps, intelligence analysts, law enforcement officials, and others is striking. As importantly, because cyberconflict is interdisciplinary, career

paths and opportunities for specialists in this area are few in number. Accordingly, the committee believes that the U.S. government should make significant efforts to develop human capital with expertise in the issues related to cyberattack.

Recommendation 10: The U.S. government should consider the establishment of a government-based institutional structure through which selected private sector entities can seek immediate relief if they are the victims of cyberattack.

As suggested in Finding 7, the United States lacks mechanisms for responding effectively to prevent further harm if a private sector entity is subjected to a cyberattack.

Given the numerous cyberattacks endured by U.S. private sector entities, it would not be surprising if one or more of these entities have taken self-help action in the past. And it is further likely that in the absence of meaningful and effective mechanisms to prevent further damage in the wake of a cyberattack, some such parties will seriously contemplate taking such action in the future if they feel that the costs of such action are less than the benefits from neutralizing the incoming attack, even if such actions constitute a violation of the Computer Fraud and Abuse Act (Section 5.2).

The argumentation for Finding 7 noted some of the undesirable aspects of taking self-help action. But the committee does not believe that a simple prohibition on such action, or even raising the penalties for such action, are alone sufficient to prevent all self-help actions in the future. For this reason, it may be desirable to consider the establishment of a government-regulated institutional structure through which private sector entities that are the targets of sustained and ongoing cyberattack can seek immediate relief.

A boundary condition in determining the appropriate structure is the impact of similar developments in other nations. That is, the U.S. government should consider the impact on the United States if other nations were to develop similar institutional structures to protect their own private sector entities.

In the absence of further study, the committee makes no endorsement of specific elements that should be included in the structure proposed in Recommendation 10. The following elements are listed for illustrative purposes only, and it should be noted that committee members disagreed among themselves about the desirability of some of these as elements of a structure for helping private sector victims of a cyberattack.

- *Improvements in capabilities for threat warning and attack assessment to*

support better forensics. Such improvements are a necessary precondition if active threat neutralization is to be a viable policy option.

- *International agreements that bind signatories to respond quickly with law enforcement actions to suppress cyberattacks emanating from their territory, with failure to do so entitling the target of the cyberattack to seek threat neutralization in response if it is located in a signatory nation.*

- *An explicit clarification of the limits to defense of property for violating the Computer Fraud and Abuse Act, which could explicitly allow or prohibit cyberattacks for this purpose.*

- *An explicit clarification of whether the victim of a cyberattack is permitted to non-destructively gather intelligence on the attacker in a non-cooperative manner. If allowed, such activities would have to be documented meticulously to demonstrate the lack of hostile intent.*

- *A capability for gathering the information needed to effect threat neutralization, accompanied by explicit rules and regulation, perhaps established by statute, to specify:*

- The selected private sector entities that are entitled to call on the government to exercise this capability for threat neutralization and the standards of security practice required of such entities;²⁹

- The circumstances under which threat neutralization is to be performed;

- The criteria needed to identify the attacking party with sufficiently high confidence; and

- The evidence needed to make the determination that any given cyberattack posed a threat sufficiently severe to warrant neutralization.

Again, to be clear, the committee does not recommend that any specific element in the list above be included or excluded in the institutional structure proposed for consideration in Recommendation 10. For example, some committee members believe that a government capability for threat neutralization is a necessary element of a robust deterrence posture against cyberattack on private sector entities, and they argue that entities under attack should themselves be allowed to effect threat neutralization subject to appropriate government regulation. Others believe it would be a serious mistake to erode the government's legal monopoly on cyber violence, and that such a capability, even if invoked promptly, would have

²⁹ The term "selected" is used in recognition of the fact that not all such entities necessarily warrant access to the institutional structure considered in Recommendation 10, and thus some mechanism will be necessary for selecting those entities that are deemed eligible. "Standards of security practice" refers to the fact that these entities should be required to adhere to good security practices as a necessary prior condition before calling for outside assistance.

at best a minimal impact in providing relief to the private sector entities under attack. Despite such disagreements, the committee does believe that it is important for the U.S. government to consider what can be done to help private sector entities cope with the undeniable inadequacies of passive defense as things currently stand.

1.9.4 Developing New Knowledge and Insight into a New Domain of Conflict

Recommendation 11: The U.S. government should conduct high-level wargaming exercises to understand the dynamics and potential consequences of cyberconflict.

As noted in Chapter 9, the dynamics of cyberconflict are not well understood, and many of the most interesting questions about cyberconflict concern matters related to deterrence, compulsion, and escalation. What are the elements that contribute to stability when cyberconflict is possible? What causes cyber adversaries to be deterred from taking hostile action? How might cyberwarfare escalate? Significant insight into crisis stability, deterrence, escalation, and other issues related to cyberconflict might be gained by conducting serious high-level wargaming exercises involving individuals with policy backgrounds and others with operational experience in cyberattack. The participation of active-duty and in-office individuals would also help to familiarize them with some of the issues. As importantly, a “gamemaster” with detailed technical knowledge of cyberdefenses and what is and is not possible through cyberattack would be essential for such exercises to produce useful knowledge. The insight and knowledge gained would be useful to senior decision makers (who would become more familiar with the issues involved), to analysts (who would gain insight into how decision makers think about such issues), and to operational personnel—the warfighters—who would gain experience in the same way that regular exercises help traditional forces develop expertise.

Recommendation 12: Foundations and government research funders should support academic and think-tank inquiry into cyberconflict, just as they have supported similar work on issues related to nuclear, biological, and chemical weapons.

The committee believes that cyberconflict and cyberattack are topics that are both important and understudied. Much of the serious thought about such subjects to date has originated in the Department of Defense, and much of that work has been classified. Whether or not the commit-

tee's recommendation is adopted regarding declassification of the policy-related discussion of cyberattack, the nation can only be better served by more open debate, discourse, and scholarship across the intellectual spectrum.

As noted in the Preface to this report, a greater interest in and more open intellectual activity regarding the subject of cyberattack would constitute an important mark of success for this committee's efforts.

Some important technical issues worth investigation include the following:

- *Attribution of cyberattacks.* Arguably the most salient technical issue in cyberconflict, other reports have underscored both the importance and the difficulty of solving the attribution problem.³⁰ This report emphatically reiterates those conclusions.

- *Attack identification.* Knowing that a nation or even a particular facility is under serious cyberattack is highly problematic given the background noise of ongoing cyberattacks all the time.

- *Geolocation of a computer that might be remotely attacked.* Given that computers are physical objects, any computer that might be attacked is in some physical location. Knowledge of that location may be important in understanding the political impact of any given cyberattack.

- *Techniques for limiting the scope of a cyberattack.* Associated with a kinetic munition is the notion of a lethal radius outside of which a given type of target is likely to be relatively unharmed. Lethal radius is a key construct for minimizing collateral damage when such munitions are used. In a world of interconnected computers, what might be a plausible analog for a "lethal radius" for cyberweapons?

There are also a host of non-technical issues raised by some of the discussion in this report. For example:

- How might cyberattack best be used to undermine the confidence of users in their information technology systems? What are the characteristics of the minimum attack needed to achieve this goal?

- What might be the impact on conflict escalation of inhibiting cyber offensive actions early in a tense international situation?

- How might cyberattack be used to support information operations such as propaganda?

- What are the relative advantages and disadvantages of different declaratory policies regarding cyberattack?

³⁰ National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington D.C., 2007.

- What are the relative advantages and disadvantages of different policies regarding self-help actions by private sector entities that come under cyberattack themselves?
- What are the dynamics of known instances of cyberattack and cyberconflict? How did the parties learn they were under attack? How did they decide to respond? What were the ramifications of responding?

1.10 CONCLUSION

Cyberattack technologies bring to the forefront of policy a wide range of possibilities in many dimensions: They raise many new policy issues, they provide many more options for operational commanders, and they complicate existing legal regimes to a considerable extent. But the findings of this report illustrate that thinking about U.S. acquisition and use of cyberattack capabilities need not start from scratch. Although a number of important nuances and subtleties can significantly complicate policy making regarding cyberattack, cyberattack should not be regarded as a *sui generis* form of warfare, and there is much to be said for drawing analogies to existing procedures, practices, and intellectual paradigms. At the same time, developing new knowledge is likely to be essential for genuinely informed policy making regarding cyberattack.

The thinking of the U.S. government on the topic of cyberattack is changing rapidly even as this report is being written. Because most of this ferment takes place behind the shields of classification, it is impossible to provide in an unclassified study a definitive report on what is going on today within the U.S. government, and it is entirely possible that some of the findings articulated and discussed above are already reflected in parts of the U.S. government and that some of the recommendations are already being implemented. If so, the committee applauds such actions. But for those findings and recommendations that have not been incorporated into government processes and thinking, the committee hopes that they will be seriously considered and that they will stimulate a government reexamination of its thinking in the relevant areas.

Part I

Framing and Basic Technology

Part I contains one chapter—Chapter 2—which provides an introduction to the technological and operational dimensions of cyberattack. The technological dimensions refer to what cyberattacks are and how they might be conducted. As the chapter makes clear, there are many different kinds of cyberattack with many different kinds of objectives, and the term “cyberattack” without further qualification should be seen more as a statement about the use of a particular attack methodology than about its targets or purpose. The operational dimensions refer to the support that a successful cyberattack requires, such as intelligence information about its targets and ways to start, stop, and calibrate a cyberattack. Cyberexploitation is addressed separately and in contrast to cyberattack.

Note to the reader: When the name of a nation is needed in this report, the names “Zendia” and “Ruritania” are used as stand-ins. Depending on context, these nations may be a near-peer nation-state with military and economic stature and power comparable to that of the United States; a small, relatively undeveloped nation; or something in between. Generally in this report, Zendia is an adversary of the United States.

Technical and Operational Considerations in Cyberattack and Cyberexploitation

This chapter focuses on technical and operational dimensions of cyberattack and cyberexploitation. Section 2.1 provides the essential points of the entire chapter, with the remainder of the chapter providing analytical backup. Section 2.2 addresses the basic technology of cyberattack. Section 2.3 addresses various operational considerations associated with “weaponizing” the basic technology of cyberattack. These sections are relevant both to the attacker, who uses cyberattack as a tool of his own choosing, and to the defender, who must cope with and respond to incoming cyberattacks launched by an attacker. Section 2.4 focuses on the centrally important issue of characterizing an incoming cyberattack. Cyberattack and cyberdefense are sometimes intimately related through the practice of active defense (Section 2.5), which may call for the defender to launch a cyberattack itself in response to an incoming cyberattack on it. Section 2.6 addresses cyberexploitation and how its technical and operational dimensions differ from cyberattack. Section 2.7 provides some lessons that can be learned from examining criminal use of cyberattack and cyberexploitation.

For perspective on tools used for cyberattack, Table 2.1 provides a comparison of tools for kinetic attack and tools for cyberattack.

Note: The committee has no specific information on actual U.S. cyberattack or cyberexploitation capabilities, and all references in this chapter to U.S. cyberattack or cyberexploitation capabilities are entirely hypothetical, provided for illustrative purposes only.

TABLE 2.1 A Comparison of Key Characteristics of Cyberattack Versus Kinetic Attack

	Kinetic Attack	Cyberattack
Effects of significance	Direct effects usually more important than indirect effects	Indirect effects usually more important than direct effects
Reversibility of direct effects	Low, entails reconstruction or rebuilding that may be time-consuming	Often highly reversible on a short time scale
Acquisition cost for weapons	Largely in procurement	Largely in research and development
Availability of base technologies	Restricted in many cases	Widespread in most cases
Intelligence requirements for successful use	Usually smaller than those required for cyberattack	Usually high compared to kinetic weapons
Uncertainties in planning	Usually smaller than those involved in cyberattack	Usually high compared to kinetic weapons

2.1 IMPORTANT CHARACTERISTICS OF CYBERATTACK AND CYBEREXPLOITATION

For purposes of this report, cyberattack refers to the use of deliberate actions—perhaps over an extended period of time—to alter, disrupt, deceive, degrade, or destroy adversary computer systems or networks or the information and/or programs resident in or transiting these systems or networks. Several characteristics of weapons for cyberattack are worthy of note:

- The indirect effects of such weapons are almost always more consequential than the direct effects of the attack. (Direct or immediate effects are effects on the computer system or network attacked. Indirect or follow-on effects are effects on the systems and/or devices that the attacked computer system or network controls or interacts with, or on the people that use or rely on the attacked computer system or network.) That is, the computer or network attacked is much less relevant than the systems controlled by the targeted computer or network or the decision making that depends on the information contained in or processed by the targeted computer or network, and indeed the indirect effect is often the primary purpose of the attack. Furthermore, the scale of damage of a cyberattack can span an enormous range.

- The outcomes of a cyberattack are often highly uncertain. Minute details of configuration can affect the outcome of a cyberattack, and cascading effects often cannot be reliably predicted. One consequence can be that collateral damage and damage assessment of a cyberattack may be very difficult to estimate.

- Cyberattacks are often very complex to plan and execute. They can involve a much larger range of options than most traditional military operations, and because they are fundamentally about an attack's secondary and tertiary effects, there are many more possible outcome paths whose analysis often requires highly specialized knowledge. The time scales on which cyberattacks operate can range from tenths of a second to years, and the spatial scales may be anywhere from "concentrated in a facility next door" to globally dispersed.

- Compared to traditional military operations, cyberattacks are relatively inexpensive. The underlying technology for carrying out cyberattacks is widely available, inexpensive, and easy to obtain. An attacker can compromise computers belonging to otherwise uninvolved parties to take part in an attack activity; use automation to increase the amount of damage that can be done per person attacking, increase the speed at which the damage is done, and decrease the required knowledge and skill level of the operator of the system; and even steal the financial assets of an adversary to use for its own ends. On the other hand, some cyberattack weapons are usable only once or a few times.

- The identity of the originating party behind a significant cyberattack can be concealed with relative ease, compared to that of a significant kinetic attack. Cyberattacks are thus easy to conduct with plausible deniability—indeed, most cyberattacks are inherently deniable. Cyberattacks are thus also well suited for being instruments of catalytic conflict—instigating conflict between two other parties.

Cyberexploitations are different from cyberattacks primarily in their objectives and in the legal constructs surrounding them. Yet, much of the technology underlying cyberexploitation is similar to that of cyberattack, and the same is true for some of the operational considerations as well. A successful cyberattack requires a vulnerability, access to that vulnerability, and a payload to be executed. A cyberexploitation requires the same three things—and the only difference is in the payload to be executed. That is, what technically distinguishes a cyberexploitation from a cyberattack is the nature of the payload. These technical similarities often mean that a targeted party may not be able to distinguish easily between a cyberexploitation and a cyberattack—a fact that may result in that party's making incorrect or misinformed decisions. On the other hand, the primary technical requirement of a cyberexploitation is that the delivery and execution

of its payload must be accomplished quietly and undetectably—secrecy is often far less important when cyberattack is the mission.

2.2 THE BASIC TECHNOLOGY OF CYBERATTACK¹

Perhaps the most important point about cyberattack from the standpoint of a major nation-state, backed by large resources, national intelligence capabilities, and political influence is that its cyberattack capabilities dwarf the kinds of cyberattacks that most citizens have experienced in everyday life or read about in the newspapers. To use a sports metaphor, the cyberattacks of the misguided teenager—even sophisticated ones—could be compared to the game that a good high school football team can play, whereas the cyberattacks that could be conducted by a major nation-state would be more comparable to the game of a professional football team with a 14-2 win-loss record in the regular season.

2.2.1 Information Technology and Infrastructure

Before considering the basic technology of cyberattack, it is helpful to review a few facts about information technology (IT) and today's IT infrastructure.

- The technology substrate of today's computers, networks, operating systems, and applications is not restricted to the U.S. military, or even just to the United States. Indeed, it is widely available around the world, to nations large and small, to subnational groups, and even to individuals.
- The essential operating parameters of this technology substrate are determined largely by commercial needs rather than military needs. Military IT draws heavily on commercial IT rather than the reverse.
- A great deal of the IT infrastructure is shared among nations and between civilian and military sectors, though the extent of such sharing varies by nation. Systems and networks used by many nations are built by the same IT vendors. Government and military users often use commercial Internet service providers. Consequently, these nominally private entities exert considerable influence over the environment in which any possible cyberconflict might take place.

¹ A primer on cyberattack in a military context can be found in Gregory Rattray, *Strategic Warfare in Cyberspace*, MIT Press, Cambridge, Mass., 2001. Rattray's treatment covers some of the same ground covered in this chapter.

2.2.2 Vulnerability, Access, and Payload

A successful cyberattack requires a vulnerability, access to that vulnerability, and a payload to be executed.² In a non-cyber context, a vulnerability might be an easily pickable lock in the file cabinet. Access would be an available path for reaching the file cabinet—and from an intruder’s perspective, access to a file cabinet located on the International Space Station would pose a very different problem from that posed by the same cabinet being located in an office in Washington, D.C. The payload is the action taken by the intruder after the lock is picked. For example, he can destroy the papers inside, or he can alter some of the information on those papers.

2.2.2.1 Vulnerabilities

For a computer or network, a vulnerability is an aspect of the system that can be used by the attacker to compromise one or more of the attributes described in the previous section. Such weaknesses may be accidentally introduced through a design or implementation flaw. They may also be introduced intentionally. An unintentionally introduced defect (“bug”) may open the door for opportunistic use of the vulnerability by an attacker who learns of its existence. Many vulnerabilities are widely publicized after they are discovered and may be used by anyone with moderate technical skills until a patch can be disseminated and installed.³ Attackers with the time and resources may also discover unintentional defects that they protect as valuable secrets—also known as zero-day exploits.⁴ As long as those defects go unaddressed, the vulnerabilities they create may be used by the attacker.

² In the lexicon of cybersecurity, “using” or “taking advantage” of a vulnerability is often called “exploiting a vulnerability.” Recall that Chapter 1 uses the term “cyberexploitation” in an espionage context—a cyber offensive action conducted for the purpose of obtaining information. The context of usage will usually make clear which of these meanings of “exploit” is intended.

³ The lag time between dissemination of a security fix to the public and its installation on a specific computer system may be considerable, and it is not always due to unawareness on the part of the system administrator. It sometimes happens that the installation of a fix will cause an application running on the system to cease working, and administrators may have to weigh the potential benefit of installing a security fix against the potential cost of rendering a critical application non-functional. Attackers take advantage of this lag time to exploit vulnerabilities.

⁴ A zero-day attack is a previously unseen attack on a previously unknown vulnerability. The term refers to the fact that the vulnerability has been known to the defender for zero days. (The attacker has usually known of the attack for a much longer time.) The most dangerous is a zero-day attack on a remotely accessible service that runs by default on all versions of a widely used operating system distribution. These types of remotely accessible

Two additional factors have increased opportunities for the attacker. First, the use of software in society has grown rapidly in recent years, and the sheer amount of software in use continues to expand across societal functions. For instance, a study by the Center for Strategic and International Studies estimated that the amount of software used in Department of Defense systems has been increasing rapidly with no let-up for the foreseeable future.⁵ More software in use inevitably means more vulnerabilities.

Second, software has also grown in complexity. Users demand more and more from software, and thus the complexity of software to meet user requirements increases steadily. Complex software, in turn, is difficult to understand, evaluate, and test.⁶ In addition, software is generally developed to provide functionality for a wide range of users, and for any particular user only a limited set of functionality may actually be useful. But whether used or not, every available capability presents an opportunity for new vulnerabilities. Simply put, unneeded capability means unnecessary vulnerability.⁷ Even custom systems often include non-essential but “nice-to-have” features that from a security perspective represent added potential for risk, and the software acquisition process is often biased in favor of excess functionality (seen as added value) while failing to properly evaluate added risk.

Of course, vulnerabilities are of no use to an attacker unless the attacker knows they are present on the system or network being attacked. But an attacker may have some special way of finding vulnerabilities, and nation-states in particular often have special advantages in doing so. For example, although proprietary software producers jealously protect their source code as intellectual property upon which their business is dependent, some such producers are known to provide source-code access to governments under certain conditions.⁸

zero-day attacks on services appear to be less frequently found as time goes on. In response, a shift in focus to the client side has occurred, resulting in many recent zero-day attacks on client-side applications. For data and analysis of zero-day attack trends, see pages 278-287 in Daniel Geer, *Measuring Security*, Cambridge, Mass., 2006, available at <http://geer.tinhow.net/measuringsecurity.tutorialv2.pdf>.

⁵ Center for Strategic and International Studies, “An Assessment of the National Security Software Industrial Base,” presented at the National Defense Industrial Association Defense Software Strategy Summit, October 19, 2006, available at http://www.diig-csis.org/pdf/Chao_SoftwareIndustrialBase_NDIASoftware.pdf.

⁶ Defense Science Board, “Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software,” U.S. Department of Defense, September 2007, p. 19.

⁷ Defense Science Board, “Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software,” U.S. Department of Defense, September 2007, p. 55.

⁸ See, for example, <http://www.microsoft.com/industry/publicsector/government/programs/GSP.msp>.

Availability of source code for inspection increases the likelihood that the inspecting party (government) will be able to identify vulnerabilities not known to the general public. Furthermore, through covert and non-public channels, nation-states may even be able to persuade vendors or willing employees of those vendors to insert vulnerabilities—secret “back doors”—into commercially available products (or require such insertion as a condition of export approval), by appealing to their patriotism or ideology, bribing or blackmailing or extorting them, or applying political pressure.

In other situations, a nation-state may have the resources to obtain (steal, buy) an example of the system of interest (perhaps already embedded in a weapons platform, for example). By whatever means the system makes its way into the hands of the nation-state, the state has the resources to test it extensively to understand its operational strengths and weaknesses, and/or conduct reverse engineering on it to understand its various functions and at least some of its vulnerabilities.

Some of the vulnerabilities useful to cyberattackers include the following:

- *Software.* Application or system software may have accidentally or deliberately introduced flaws whose use can subvert the intended purpose for which the software is designed.
- *Hardware.* Vulnerabilities can also be found in hardware, including microprocessors, microcontrollers, circuit boards, power supplies, peripherals such as printers or scanners, storage devices, and communications equipment such as network cards. Tampering with such components may secretly alter the intended functionality of the component, or provide opportunities to introduce hostile software.
- *Seams between hardware and software.* An example of such a seam might be the reprogrammable read-only memory of a computer (firmware) that can be improperly and clandestinely reprogrammed.
- *Communications channels.* The communications channels between a system or network and the “outside” world can be used by an attacker in many ways. An attacker can pretend to be an “authorized” user of the channel, jam it and thus deny its use to the adversary, or eavesdrop on it to obtain information intended by the adversary to be confidential.
- *Configuration.* Most systems provide a variety of configuration options that users can set, based on their own security versus convenience tradeoffs. Because convenience is often valued more than security, many systems are—in practice—configured insecurely.
- *Users and operators.* Authorized users and operators of a system or network can be tricked or blackmailed into doing the bidding of an attacker.

- *Service providers.* Many computer installations rely on outside parties to provide computer-related services, such as maintenance or Internet service. An attacker may be able to persuade a service provider to take some special action on its behalf, such as installing attack software on a target computer.

Appendix E discusses these vulnerabilities in more detail.

2.2.2.2 Access

In order to take advantage of a vulnerability, a cyberattacker must have access to it. Targets that are “easy” to attack are those that involve relatively little preparation on the part of the attacker and where access to the target can be gained without much difficulty—such as a target that is known to be connected to the Internet. Public websites are a canonical example of such targets, as they usually run on generic server software and are connected to the Internet, and indeed website defacement is an example of a popular cyberattack that can be launched by relatively unskilled individuals.

At the other end of the spectrum, difficult targets are those that require a great deal of preparation on the part of the attacker and where access to the target can be gained only at great effort or may even be impossible for all practical purposes. For example, the on-board avionics of an adversary’s fighter plane are not likely to be connected to the Internet for the foreseeable future, which means that launching a cyberattack against it will require some kind of close access to introduce a vulnerability that can be used later (close-access attacks are discussed in Section 2.2.5.2). Nor are these avionics likely to be running on a commercial operating system such as Windows, which means that information on the vulnerabilities of the avionics software will probably have to be found by obtaining a clandestine copy of it. In general, it would be expected that an adversary’s important and sensitive computer systems or networks would fall into the category of difficult targets.⁹

Access paths to a target may be transient. For example, antiradiation missiles often home in on the emissions of adversary radar systems; once

⁹ An important caveat is the fact that adversary computer systems and networks are subject to the same cost pressures as U.S. systems and networks, and there is no reason to suppose that adversaries are any better at avoiding dumb mistakes than the United States is. Thus, it would not be entirely surprising to see important and/or sensitive systems connected to the Internet because the Internet provides a convenient communications medium, or for such systems to be built on commercial operating systems with known vulnerabilities because doing so would reduce the cost of development. However, the point is that no cyberattacker can *count on* such dumb mistakes for any particular target of interest.

the radar shuts down, the missile aims at the last known position of the radar. Counterbattery systems locate adversary artillery by backtracing the trajectory of artillery shells, but moving the artillery piece quickly makes it relatively untargetable. Similar considerations sometimes apply to an adversary computer that makes itself known by transmitting (e.g., conducting an attack). Under such circumstances, a successful cyberattack on the adversary computer may require speed to establish an access path and use a vulnerability before the computer goes dark and makes establishing a path difficult or impossible.

Under some other circumstances, an access path may be intermittent. For example, a submarine's onboard administrative local area network would necessarily be disconnected from the Internet while underwater at sea, but might be connected to the Internet while in port. If the administrative network is ever connected to the on-board operational network (controlling weapons and propulsion) at sea, an effective access path may be present for an attacker.

Access paths to a target can suggest a way of differentiating between two categories of cyberattack:

- *Remote-access cyberattacks*, in which an attack is launched at some distance from the adversary computer or network of interest. The canonical example of a remote access attack is that of an adversary computer attacked through the access path provided by the Internet, but other examples might include accessing an adversary computer through a dial-up modem attached to it or through penetration of the wireless network to which it is connected and then proceeding to destroy data on it.¹⁰

- *Close-access cyberattacks*, in which an attack on an adversary computer or network takes place through the local installation of hardware or software functionality by friendly parties (e.g., covert agents, vendors) in close proximity to the computer or network of interest. Close access is a possibility anywhere in the supply chain of a system that will be deployed, and it may well be easier to gain access to the system before it is deployed.

These two categories of cyberattack may overlap to a certain extent. For example, a close-access cyberattack might result in the implantation of friendly code in online, Internet-propagated updates to a widely used

¹⁰ The Department of Defense (DOD) definition of computer network attack (CNA)—“actions taken through the use of computer networks to disrupt, deny, degrade, or destroy information resident in computers and computer networks, or the computers and networks themselves”—is similar in spirit to this report's use of “remote-access” cyberattack. See Joint Publication 3-13, *Information Operations*, February 13, 2006.

program. Such an attack would embody elements of the two categories. Also, communications channels (the channels through which IT systems and networks transfer information) can also be targeted through remote access (e.g., penetrating or jamming a wireless network) or through close access (e.g., tapping into a physical cable feeding a network).

2.2.2.3 Payload

Payload is the term used to describe the things that can be done once a vulnerability has been exploited. For example, once a software agent (such as a virus) has entered a given computer, it can be programmed to do many things—reproducing and retransmitting itself, destroying files on the system, or altering files.

Payloads can have multiple capabilities when inserted into an adversary system or network—that is, they can be programmed to do more than one thing. The timing of these actions can also be varied. And if a communications channel to the attacker is available, payloads can be remotely updated. Indeed, in some cases, the initial delivered payload consists of nothing more than a mechanism for scanning the system to determine its technical characteristics and an update mechanism to retrieve from the attacker the best packages to further its attack.

A hostile payload may be a Trojan horse—a program that appears to be innocuous but in fact has a hostile function that is triggered immediately or when some condition is met. It may also be a rootkit—a program that is hidden from the operating system or virus checking software but that nonetheless has access to some or all of the computer's functions. Rootkits can be installed in the boot-up software of a computer, and even in the BIOS ROM hardware that initially controls the boot-up sequence. (Rootkits installed in this latter manner will remain even when the user erases the entire hard disk and reinstalls the operating system from scratch.)

Once introduced into a targeted system, the payload sits quietly and does nothing harmful most of the time. However, at the right moment, the program activates itself and proceeds to (for example) destroy or corrupt data, disable system defenses, or introduce false message traffic. The “right moment” can be triggered because a certain date and time are reached, because the payload receives an explicit instruction to activate through some covert channel, because the traffic it monitors signals the right moment, or because something specific happens in its immediate environment.

An example is a payload that searches for “packets of death.” This payload examines incoming data packets on a host for a special pattern embedded within it. For almost all packets, the payload does nothing. But when it sees a particular sequence of specially configured packets,

it triggers some other hostile action—it crashes the system, deletes files, corrupts subsequent data packets, and so on. (Note that the hostile action may be to do nothing when action should be taken—an air-defense system that ignores the signature of certain aircraft when it receives such a packet has clearly been compromised.)

Note that payloads for cyberattack may be selective or indiscriminate in their targeting. That is, some payloads for cyberattack can be configured to attack any computer to which access may be gained, and others can be configured to attack quite selectively only certain computers.

2.2.3 Scale and Precision

A cyberattack can be conducted over a wide range of scales, depending on the needs of the attacker. An attack intended to degrade confidence in the IT infrastructure of a nation might be directed against every Internet-connected desktop computer that uses a particular operating system. Attacks intended to “zombify” computers for later use in a botnet need not succeed against any particular machine, but instead rely on the fact that a large fraction of the machines connected to the Internet will be vulnerable to being compromised.

Alternatively, a cyberattack might be directed to all available targets across one or more critical infrastructure sectors. A probe intended to test the feasibility of a large-scale cyberattack might be directed against just a few such computers selected at random. An attack might also be directed against a few selected key targets in order to have secondary effects (e.g., disruption of emergency call dispatch centers timed to coincide with physical attacks, thus amplifying the psychological effect of those physical attacks).

A cyberattacker may also care about which computers or networks are targeted—an issue of precision. Of greatest significance are the scenarios in which focused but small-scale attacks are directed against a specific computer or user whose individual compromise would have enormous value (“going after the crown jewels”)—an adversary’s nuclear command and control system, for example. Or, a cyberattack may be directed against a particular electric power generation plant that powers a specific building in which adversary command and control systems are known to operate, rather than all of the generation facilities in a nation’s entire electric grid.

2.2.4 Critical Periods of Cyberattack

How a cyberattack evolves over time is relevant, and there are several time periods of interest. The first, $T_{\text{intelligence collection}}$, is the period

available for collecting intelligence needed to launch the attack. A second relevant period, $T_{\text{attack launch}}$, is the period over which the functionality required to carry out the attack on the targeted system(s) is installed or deployed—that is, during which the attack is launched. A third relevant period, $T_{\text{compromise}}$, is the period over which the confidentiality, integrity, or availability attributes of the targeted system(s) are compromised. A fourth relevant period, $T_{\text{effects apparent}}$, is the time period over which the victim actually suffers the ill effects of such compromises. During this time, the target can recover from the attack or reconstitute its function. Depending on the specific nature of the cyberattack, these four periods may—or may not—overlap with each other.

The distinctions between these various periods are important.¹¹ For example, the fact that $T_{\text{attack launch}}$ and $T_{\text{compromise}}$ are different windows in time means that the period $T_{\text{attack launch}}$ can be used to “pre-position” vulnerabilities to facilitate later actions. This pre-positioning could be in the form of trapdoors left behind from previous virus infections, unintentional design vulnerabilities,¹² or vulnerable code left by a compromised staff member or by a break-in to the developer’s site.¹³

Such pre-positioning is helpful for launching high-volume cyberattacks—possible targets include air-traffic control facilities, systems in manufacturing or shipping departments, logistics systems in rail transport companies, energy production facilities, and so on, as well as a variety of military facilities. An attacker that has prepared his targets in this manner has avenues for instantaneous disruption or corruption of operational processes through a large-scale injection of forged communications, destruction of data, or flooding of services from inside normal perimeter defenses. When hosts inside a network begin to attack the internal network infrastructure or servers, they are often hard to identify rapidly because the very tools that are used by network operations staff to diagnose network problems may not be available. An attack spread widely enough can overwhelm the network operations and system

¹¹ These concepts can also be found in epidemiologic models for the spread of malware. See, for example, <http://geer.tinho.net/measuringsecurity.tutorialv2.pdf>.

¹² An example is the recent episode during which Sony’s BMG Music Entertainment surreptitiously distributed software on audio compact disks (CD) that was automatically installed on any computers that played the CDs. This software was intended to block the copying of the CD, but it had the unintentional side effect of opening security vulnerabilities that could be exploited by other malicious software such as worms or viruses. See Iain Thomson and Tom Sanders, “Virus Writers Exploit Sony DRM,” November 10, 2005, available at <http://www.vnunet.com/vnunet/news/2145874/virus-writers-exploit-sony-drm>.

¹³ P.A. Karger and R.R. Schell, *Multics Security Evaluation: Vulnerability Analysis*, ESD-TR-74-193, Vol. II, June 1974, HQ Electronic Systems Division, Hanscom Air Force Base, available at <http://csrc.nist.gov/publications/history/karg74.pdf>.

administration staffs, increasing the time it takes to diagnose and mitigate the attack.

More complex attacks can also be coordinated with other events to achieve a force-multiplier effect. For example, if an attack of this nature could successfully be made against an air defense network, the attacker could disrupt the network's operation in concert with a hostile flight operation, potentially blinding the defense system for a period of time.

Still another relevant time scale is the duration of what might be called the entire operation, which itself might call for multiple cyberattacks to be conducted over time. A denial-of-service attack is the canonical example of an operation that requires multiple cyberattacks over a period of time—when the attacks stop, the denial of service vanishes. Multiple cyberattacks conducted over time might also be needed to coordinate the entire cyber operation with other military activity taken against an adversary. Alternatively, and perhaps more prosaically, multiple cyberattacks might be needed to ensure the continuing disruption of an adversary computer system or network because the vulnerabilities that an attacker needs to target may not remain static. In an operation that calls for multiple cyberattacks over time, the targeted party may well respond to the first signs of the attack by closing or correcting some or all of the vulnerabilities enabling the attack. Other vulnerabilities would no doubt remain, but the attacker would have to have advance knowledge of them and adjust the attack accordingly.

These different time scales also help to explain possible different perceptions of the parties involved regarding what might “count” as an attack. For example, the attacker might reasonably believe that a cyberattack has not been committed until a hostile agent planted on the adversary's computer has caused actual damage to it. The adversary might reasonably believe that a cyberattack has been committed when the agent was first planted on the computer, thus giving the agent's controller the technical *capability* of causing actual damage to the adversary's computer.

2.2.5 Approaches for Cyberattack

What are the approaches that may be used for cyberattack? In general, a cyberattack depends on the attacker's operational skill and knowledge of the adversary, and its success relies on taking advantage of a mix of the adversary's technical and human vulnerabilities. Furthermore, although the military services and intelligence agencies rightly classify a variety of operational techniques and approaches to cyberattack, they too are governed by the same laws of physics as everyone else, and there are no magic technologies behind closed doors that enable, for example, an

attacker to beam a virus into a computer that lacks connections to the outside world.

The discussion in this section is thus based on what is publicly known about tools and methods for conducting cyberattacks. This generic discussion is not intended to be complete or comprehensive, but it provides enough information for readers to gain a sense of what might be possible without providing a detailed road map for a would-be cyberattacker. This section is divided into three subsections: malware suitable for remote attacks, approaches for close-access attacks, and social engineering.

In many cases, these tools and methods are known because they have been used in criminal enterprises. Box 2.1 describes some of the advantages that a nation-state has over others in conducting cyberattacks; these advantages can be used to refine and weaponize these tools and methods so that they are more effective in carrying out operational missions. (The actions of nation-states are also constrained by applicable domestic laws, although in the United States, certain government agencies are explicitly exempted from complying with certain laws. See Section 7.3.4 for more discussion of this point.)

2.2.5.1 Possible Approaches for Remote-Access Cyberattacks

Remote-access cyberattacks can be facilitated with a variety of tools, some of which are described below.

2.2.5.1.1 Botnets

An attack technology of particular power and significance is the botnet. Botnets are arrays of compromised computers that are remotely controlled by the attacker. A compromised computer—an individual bot—is connected to the Internet, usually with an “always-on” broadband connection, and is running software clandestinely introduced by the attacker. The attack value of a botnet arises from the sheer number of computers that an attacker can control—often tens or hundreds of thousands and perhaps as many as a million. (An individual unprotected computer may be part of multiple botnets as the result of multiple compromises.) Since all of these computers are under one party’s control, the botnet can act as a powerful amplifier of an attacker’s actions.

An attacker usually builds a botnet by finding a few individual computers to compromise, perhaps using one of the tools described above. The first hostile action that these initial zombies take is to find other machines to compromise—a task that can be undertaken in an automatic manner, and so the size of the botnet can grow quite rapidly. It is widely reported that only minutes elapse between the instant that a computer

BOX 2.1 Cyberattack Advantages of Nation-States over Other Types of Actors

Nations have enormous resources to bring to bear on the problem of cyber-attack. They are financed by national treasuries; they can take advantage of the talents of some of the smartest and most motivated individuals in their populations; they often have the luxury of time to plan and execute attacks; and they can draw on all of the other resources available to the national government, such as national intelligence, military, and law enforcement services. As a result, a government-supported cyberattacker can be relatively profligate in executing its attack and in particular can target vulnerabilities at any point in the information technology supply chain from hardware fabrication to user actions.

The availability of such resources widens the possible target set of nation-state attackers. Low- and mid-level attackers often benefit from the ability to gain a small profit from each of many targets. Spammers and “bot” harvesters are the best examples of this phenomenon—an individual user or computer is vulnerable in some way to a spammer or a bot harvester, but the spammer or bot harvester profits because many such users or computers are present on the Internet. However, because of the resources available to them, high-end attackers may also be able to target a specific computer or user whose individual compromise would have enormous value (“going after the crown jewels”). In the former case, an attacker confronted with an adequately defended system simply moves on to another system that is not so well defended. In the latter case, the attacker has the resources to increase the scale and sophistication of the attack to a very high degree if the target is sufficiently valuable.

It is also the case that the resources available to a nation are not static. This means that for a sufficiently valuable target, a nation may well be able to deploy additional resources in its continuing attack if its initial attacks fail. In other words, capabilities that are infeasible for a nation today may become feasible tomorrow.

A nation-state also has the resources that allow it to obtain detailed information about the target system, such as knowledge gained by having access to the source code of the software running on the target or the schematics of the target device or through reverse engineering. (A proof of principle is illustrated in the short delay between the unauthorized public disclosure of portions of the source code for Microsoft Windows 2000 and Windows NT 4.0 in February 2004 and the reported appearance of exploits apparently derived from an examination of the source code.¹) Success in obtaining such information is not guaranteed, of course, but the likelihood of success is clearly an increasing function of the availability of resources.

A nation-state cyberattacker does not care how it succeeds, as long as the path to success meets various constraints such as affordability and secrecy. In particular, the nation-state has the ability to compromise or blackmail a trusted insider to do its bidding or infiltrate a target organization with a trained agent rather than crack a security system if the former is easier to do than the latter.

¹ Statement from Microsoft Regarding Illegal Posting of Windows 2000 Source Code, <http://www.microsoft.com/presspass/press/2004/Feb04/02-12windowssource.msp>.

BOX 2.2 Managing a Botnet

Many botnets today make use of a two- or three-tier hierarchy, from attacker to a central controller (or handler), from handler to agents (“bots”), and sometimes from agents (bots) to reflectors.¹ There are still limitations, however, on how many bots can be in a given channel at a given time, and attackers using large botnets know to move them around from server to server and channel to channel (known as herding) to avoid discovery or take-down of the botnets.

A high-capacity attack network can make good use of another layer between the attacker and the handlers, which ideally is highly survivable and hardened so as to remain active in the face of defensive action. This layer is then used to control multiple independent botnets, or lower levels of distributed attack networks, in a manner similar to the regiment/battalion/company hierarchy used by conventional military forces. By adding this additional layer, it is possible to coordinate much larger forces using independent teams, similar to what was done by the team of “consultants” in Operation Cyberslam,² but on an even larger scale. This additional layer will require a database to keep track of the various lower-level distributed attack networks and to assemble, reassemble, or reconstitute them as need be, in a manner very similar to maintaining force size through replacement of killed or wounded soldiers, and adjusting force strength through redeployment and reinforcement.

Another approach to command and control involves two-way communications between controllers and bots. For example, rather than await orders, bots can send requests to controllers asking what to do next (“pulling” orders rather than “pushing” them). Successive requests from one bot go to different controllers. If a controller does not respond, after a while the bot tries another controller. If none of the controllers respond, the bot generates a series of random Domain Name System (DNS) names and tries those hosts, one at a time. The bot herders know

attaches to the Internet and the time that it is probed for vulnerabilities and possibly compromised itself.¹⁴

A botnet attacker (controller) can communicate with its botnet (Box 2.2) and still stay in the background, unidentified and far away from any action, while the individual bots—which may belong mostly to innocent parties that may be located anywhere in the world—are the ones that are visible to the party under attack. The botnet controller has great flexibility

¹⁴ See, for example, *Survival Time*, available at <http://isc.sans.org/survivaltime.html>. Also, in a 2008 experiment conducted in Auckland, New Zealand, an unprotected computer was rendered unusable through online attacks. The computer was probed within 30 seconds of its going online, and the first attempt at intrusion occurred within the first 2 minutes. After 100 minutes, the computer was unusable. (See “Experiment Highlights Computer Risks,” December 2, 2008, available at <http://www.stuff.co.nz/print/4778864a28.html>.)

the random number generation algorithm, and if they lose control of some group of bots (perhaps because some controllers have been discovered and disabled), they register one of the random DNS names just before the orphaned bots are about to try it, and when the bot checks in they update it to regain control.

It is also possible to use out-of-band communications (e.g., telephone, radio, or face-to-face conversation) to relay targeting information and attack timing. Especially in the case of long-running operations, there is no need for constant or immediate network connections between attacking networks. In fact, it would be less expensive and less risky from an operational security perspective to coordinate a large number of distributed attack networks independently of each other. In this way, one or more groups could be responsible for recruiting (compromising and taking control over) new computers, which are then added to individual attack networks as requested when capacity drops below a certain level. If the attack tools are designed in a sufficiently modular way—and IRC botnets today already have these capabilities built in—this becomes an issue of human management rather than technology.

¹ An example of a reflector attack would be to send DNS requests with a forged source address containing the intended target's IP address to a large number of DNS servers, which would in turn send the replies back to what they believed to be the "requester" (the victim), which is then flooded with traffic. If the DNS request packet contained 100 bytes of data, and the replies contained 700 bytes of data, a 7× amplification would result, in addition to reflection. There is no need for malicious software to be installed on the reflector; hence this makes a good indirect attack method that is very hard to trace back to the attacker.

² Department of Justice, "Criminal Complaint: United States of America v. Paul G. Ashley, Jonathan David Hall, Joshua James Schichtel, Richard Roby and Lee Graham Walker," 2004, available at <http://www.reverse.net/operationcyberslam.pdf>.

in the actions he may take—he may direct all of the bots to take the same action, or each of them to take different actions.

Botnets are ideally suited for conducting distributed denial-of-service (DDOS) attacks against computer systems, and there are few good ways to defend against such attacks. A denial-of-service attack floods a specific target with bogus requests for service, thereby exhausting the resources available to the target to handle legitimate requests for service and thus blocking others from using those resources. Such an attack is relatively easy to block if these bogus requests for service come from a single source, because the target can simply drop all service requests from that source. However, a distributed denial-of-service attack can flood the target with multiple requests from many different machines, each of which might, in principle, be a legitimate requester of service.

DDOS attacks are often conducted using unprotected machines in the

open Internet environment. But there is no reason in principle that they could not be conducted in a more closed environment, such as a classified Zendian network for military communications or planning. Planting the initial botnet “seeds” would probably be more difficult and time-consuming than doing so on the open Internet, but once established, it is likely that the “inside” botnet could grow rapidly because many sensitive networks are protected only by a hardened perimeter.

Individual bots can sense and probe their immediate environment. For example, a bot can examine clear-text data (e.g., sensitive information such as user names and passwords) passing by or through its host computer, including keystrokes entered by users and traffic on the local area network to which that host computer is attached. It might examine data files accessible to the host computer, such as any document stored on the computer. This information could be harvested and passed back to the botnet controller and mined for useful intelligence (a cyberexploitation).

A bot could examine system files on the system to ascertain the particular operating system and version being used, transmit this information back to the controller, and receive in return an upgraded payload that is specifically customized for that environment. This payload might be a destructive one, to be triggered at a certain time, or perhaps when the resident bot receives a subsequent communication from the controller. As a cyberexploitation, it could also ascertain the identity(ies) of the users and possibly their roles in an organization.

A bot could assume the identity of a legitimate user, and use its host as the originating site for e-mail. Whereas in the criminal world botnets often generate spam e-mail consisting of millions of identical messages, a military application might call for sending a personalized message from a compromised bot to another, uncompromised user that would mislead or confuse him.

Individual bots can also act as hosts for information exfiltration. Botnets sharing data using encrypted peer-to-peer connections could be used as “distributed dead drops.” This would make it much more difficult to prevent the information from being received and to discern the ultimate location of the botnet controllers.

Perhaps the most important point about botnets is the great flexibility they offer an attacker (or an exploiter). Although they are known to be well suited to DDOS attacks, it is safe to say that their full range of utility for cyberattack and cyberexploitation has not yet been examined.

2.2.5.1.2 Other Tools and Approaches for Remote-Access Cyberattack

Security Penetrations The owner or operator of an important system usually takes at least some measures to protect it against outside intruders. A

common security suite may involve requiring users to authenticate themselves and running security software that selectively blocks external access (firewalls) and checks for hostile malware that may be introduced.

However, password guessing is a common method for penetrating system security. Users have a tendency to choose easily remembered passwords that they change rarely if ever, suggesting certain patterns in password choice that are likely to be common. For example, passwords are often drawn from popular culture and sports, or are likely to be words from the user's native language. Even when the system attempts to enforce password choices with variation (e.g., "must contain a digit"), people subvert the intent in simple and easily predictable ways (e.g., they often choose PASSWORD1, PASSWORD2, and so on). Dictionaries are often used for guessing passwords—e.g., trying every word in the Zendian dictionary; such a technique can be effective if proper safeguards are not in place. Similar problems hold for any authenticator that remains constant.

An attacker may try to compromise security software to pave the way for the introduction of another attack agent. Some agents evade detection by varying the malicious payload or by checking constantly to ensure that a given virus is not identified by antivirus engines.¹⁵ Others are designed to disable antivirus programs, and may do so selectively so that only a specific virus or worm written by the attacker sent later will be allowed through. These are simple automation steps that can use techniques described openly within the computer security industry.¹⁶

Worms and Viruses Worms and viruses are techniques generally used for installing Trojan horses on many computers. A worm is self-replicating—in addition to infecting the machine on which it is resident, it uses that machine to seek out other machines to infect. A virus replicates through user actions—for example, an e-mail containing a virus may be sent to Alice. When Alice opens the e-mail, her computer is infected. The virus program may then send an e-mail (ostensibly from Alice) to every person in her contact list; when Bob receives the e-mail from Alice and opens it, Bob's computer is infected and the cycle repeats itself. Because user action is required to propagate a virus, viruses tend to spread more slowly than do worms.

Worms and viruses may be initially propagated in many ways, includ-

¹⁵ As many as 30 percent of virus and other malware infections may be undetectable by today's antivirus engines. See, for example, Niels Provos et al., "All Your iFRAMEs Point to Us," *Proceedings of the 17th Usenix Security Symposium 08*, 2008, available at http://www.usenix.org/events/sec08/tech/full_papers/provos/provos.pdf.

¹⁶ Metasploit Anti-Forensics Project. Metasploit Anti-forensics homepage, available at <http://www.metasploit.com/research/projects/antiforensics/>.

ing e-mails received by, web pages visited by, images displayed by, and software downloaded by the victim. Worms and viruses are often used as intermediate stepping stones to assume full control of an adversary system. For example, they can be used to establish reverse tunnels out through the firewall (from inside to outside), which in turn grant someone outside the protected network full control of the host inside the network, or to control hosts in an enterprise in the supply chain of the primary target.

Anonymizers Anonymizers are used to conceal the identity of an attacking party. One particularly useful anonymizing technique is onion routing, a technology originally designed to disguise the source of electronic traffic.¹⁷ But since the technology cannot distinguish between different kinds of traffic, attackers can use onion routing to disguise the source of a remote cyberattack.

Onion routing works by establishing a path through a maze of multiple onion routers, each of which accepts a packet from a previous router and forwards it on to another onion router. The originating party—in this case, the attacker—encrypts the packet multiple times in such a way that each onion router can peel off a single layer of encryption; the final router peels off the last layer, is able to read the packet in the clear, and sends it to the appropriate destination. A variety of public-domain onion router networks exist, and some support specifying where the exit point should be. Thus, the attacker can specify “Exit from an onion router located in Zendia” and that is where a target would see an attack coming from. (On the other hand, a sophisticated target might notice that an attack was coming from a public-domain onion router, and make probabilistic inferences, though not definitive, about where the attack was really coming from.)

Penetrations of and Denial-of-Service Attacks on Wireless Networks Wireless networks to enable communications among computers and devices are increasingly common and provide clandestine methods for access and denial of service. An attacker may be able to insert his own broadcast/reception node (on a WiFi network, he might insert his own wireless access point) to intercept and monitor traffic and perhaps be able to impersonate an authorized network user.

An attacker may sometimes impersonate an authorized user with relative ease if access to the wireless network is not protected. For example, satellites communicate with their ground stations through wireless communications, and the command link may not be encrypted or may

¹⁷ See, for example, the TOR router (and project) at <http://www.torproject.org/>.

be otherwise insecure. If so, a Zendian satellite can be controlled by commands sent from the United States just as easily as by commands sent from Zendia. With access to the command link, adversary satellites can be turned off, redirected, or even directed to self-destruct by operating in unsafe modes.

Alternatively, an attacker might choose to deny service on the network by jamming it, flooding the operating area with RF energy of the appropriate frequencies. WiFi wireless networks for computer communications are an obvious target, and given the increasing ubiquity of cell phones around the world, cell phone networks could be a particularly useful target of a jamming cyberattack should an attacker wish to disrupt a primary communications mechanism probably used by government and non-government personnel alike.

Router Compromises Router compromises often manipulate the logical map of a network (whether open, like the Internet, or closed, like that of a corporate network) in ways desired by the attacker. For example, modification of the software and data tables that underlie routing of information, a specific site could effectively be isolated so that it could not receive e-mail sent to it from elsewhere on the Internet or so that web pages hosted on it could never be found by anyone on the Internet. A different modification of the routing software and data might result in much more widespread chaos on the Internet for many sites rather than just one.

Attacks on routers are feasible because the routers themselves are often Internet accessible and have software and hardware vulnerabilities just like any other computers, although even if they were not Internet accessible, compromising them would not be impossible. Moreover, code to support attacks on routers is often available in the public domain, making attacks on routers easier than they would otherwise be. Under some circumstances, router flaws may enable an attacker to damage the routing hardware itself remotely, as might be possible if the boot ROM were compromised or if the attacker gained access to low-level functions that controlled power supplies or fan speeds.

An example of a router compromise is the Border Gateway Protocol (BGP) attacks. The Internet is a network of networks. Each network acts as an autonomous system under a common administration and with common routing policies. Primarily used by Internet service providers (ISPs) and very large private networks such as those of multinational corporations, BGP is the Internet protocol used to characterize every network to each other, for example between ISPs. BGP does so by publishing tables containing information on how packets should be routed between any given network and other networks. However, if these tables are cor-

rupted, traffic can be misdirected.¹⁸ One kind of BGP attack deliberately corrupts these tables to misdirect traffic away from its rightful destination and toward a network controlled by the attacker.

Once the attacker has captured traffic intended for a given destination, the captured traffic can be discarded (thus isolating the destination network) or copied for later examination and then forwarded to the correct destination (to reduce the likelihood of the attack becoming known). If the captured traffic contains information such as passwords, the attacker may be able to impersonate the sender at a later date. Another kind of attack hijacks a block of IP addresses in order to send undesirable or malicious traffic such as spam or denial-of-service attacks. Such an attack allows a sender to remain untraceable. The attacker uses the routing infrastructure to evade IP-based filtering aimed at blacklisting malicious hosts.

Protocol Compromises A network protocol is a standard that controls or enables communication between two computing devices. In practice, protocols—even widely accepted and used protocols—are sometimes flawed. For example, a given protocol may be designed incorrectly or be incompletely specified.¹⁹ A given implementation of a well-designed and well-specified protocol may itself be incomplete and/or contain a bug. (An incomplete implementation may mean that the system can enter some unanticipated state, and thus that consequences that ensue are unpredictable.) An attacker may take advantage of such flaws.

An example of a protocol attack is DNS cache poisoning. The Domain Name System (DNS) is a global system that maps domain names (e.g., *www.nas.edu*) into specific numeric IP addresses (e.g., *144.171.1.22*).²⁰ However, in order to reduce the load on the primary name servers, tables containing the relevant information are stored (cached) on secondary DNS servers operated by Internet service providers. By taking advantage

¹⁸ See Xin Hu and Z. Morley Mao, "Accurate Real-Time Identification of IP Prefix Hijacking," *Proceedings of IEEE Symposium on Security and Privacy*, May 2007, pp. 3-17; Anirudh Ramachandran and Nick Feamster, "Understanding the Network-Level Behavior of Spammers," *Proceedings of the Association of Computing Machinery SIGCOMM 2006*, pp. 291-302, available at <http://www.cc.gatech.edu/~avr/publications/p396-ramachandran-sigcomm06.pdf>.

¹⁹ Incomplete specifications or implementations are dangerous because of the possibility of inputs for which no response is specified or provided. That is, a protocol (or a given implementation of the protocol) may not unambiguously specify an action to be taken for all inputs. If one of these "undefined response" inputs is received, the receiving system will do something unanticipated. If enough is known about the receiving system and its particular implementation of the protocol being used, the subsequent action may be exploitable.

²⁰ Every device connected to the Internet has a unique identifying number known as its IP address. An IP address may take the form *144.171.1.22* (for IP Version 4) or *2001:db8:0:1234:0:567:1:1* (for IP Version 6).

of vulnerabilities in DNS software, it is sometimes possible to alter these tables, so that a request to “www.nas.edu” maps to 144.117.1.22, rather than the correct 144.171.1.22. The incorrect IP address 144.117.1.22 can be a phony host configured to look like the real thing, and can be used to intercept information sent by the user and intended for the real site. Alternatively, corrupted tables could be used simply to misdirect messages being transmitted from point to point.²¹ (The corruption of a DNS server to redirect traffic is sometimes known as “pharming.”) Cache poisoning is possible because the DNS protocol does not authenticate responses, which is widely regarded as a flaw in the security of that protocol. This flaw means that an attacker can take advantage of the protocol by sending an inquiry to a server that causes it to make an inquiry to another server, and then sending a bogus reply to the first server before the second server has a chance to respond.

Other examples of protocol attacks may involve partially opening many Transmission Control Protocol connections to tie up resources, or sending packets marked as the first and last fragments of a huge datagram in order to tie up buffer space.

2.2.5.2 Possible Approaches for Close-Access Cyberattacks

To reduce the threat from tools that enable remote attacks, a potential target might choose to disconnect from easily accessible channels of communication. A computer that is “air gapped” from the Internet is not susceptible to an attack that arrives at the computer through Internet connections. Thus, it is sometimes too difficult or impossible for an attacker to obtain remote access to a computer of interest. In these instances, other methods of attack are necessary.

One approach to attacking a putatively isolated and “stand-alone” computer is to consider whether that computer is in fact isolated. For example, a computer without an Internet connection may be accessible through a dial-up modem; if the attacker can discover the phone number associated with that modem, the computer may be vulnerable to a remote attack for the price of a long-distance telephone call. Or the computer of interest may connect to the Internet only occasionally to receive updates—during those moments of connection, the computer may be vulnerable. Or the computer might require the use of external media to provide data—although the data does not arrive through an Internet connection, data is supplied through the insertion of the data-carrying media into an appropriate slot in the computer, and the placement of hostile data

²¹ National Research Council, *Signposts in Cyberspace: The Domain Name System and Internet Navigation*, The National Academies Press, Washington, D.C., 2005.

on the CD-ROM can occur on a computer that *is* connected to the Internet. (Hostile data is data that, when processed, might cause the computer to fail or crash.)

If the computer or network of interest is indeed isolated, close-access attacks provide an alternative. Close-access attacks require a human to be physically present, which can increase the chances of being caught and/or identified. Sometimes, a close-access attack is used late in the supply chain, e.g., against a deployed computer in operation or one that is awaiting delivery on a loading dock. In these cases, the attacks are by their nature narrowly and specifically targeted, and they are also not scalable, because the number of computers that can be compromised is proportional to the number of human assets available. In other cases, a close-access cyberattack may be used early in the supply chain (e.g., introducing a vulnerability during development), and high leverage might result from such an attack. For example, for many years the United States overtly restricted the cryptographic strength of encryption products allowed for export. If it had done so covertly, such an action might well have counted as a close-access cyberattack intended to make encryption products more vulnerable to compromise.

By definition, close-access attacks bypass network perimeter defenses. A successful close-access cyberattack makes an outsider appear, for all intents and purposes, to be an insider, especially if credentials have already been compromised and can be used without raising alarms. Sophisticated anomaly detection systems that operate from login audit logs, network flow logs, and other sources of network and computer usage would be necessary to be able to detect this type of activity. Standard antivirus software and intrusion detection or protection systems are significantly less effective.

Examples of close-access cyberattacks include the following:

- Attacks somewhere in the supply chain. Systems (and their components) can be attacked in design, development, testing, production, distribution, installation, configuration, maintenance, and operation. (See Box 2.3 for a documented example.) Indeed, the supply chain is only as secure as its weakest link.²² In most cases, the supply chain is only loosely managed, which means that access control throughout the entire supply chain is difficult. Examples of hypothetical supply-chain attacks include the following:

²² Defense Science Board, "Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software," U.S. Department of Defense, September 2007, p. 25.

BOX 2.3 Project “Gunman”

On October 25, 1990, Congressman Henry Hyde revealed some of the technical highlights of a Soviet intelligence penetration of the U.S. embassy in Moscow that was still under construction at that time. Within the U.S. intelligence community, the Soviet operation was known by the code name “Gunman.” In 1984, the United States discovered that a number of IBM Selectric typewriters within the U.S. Embassy had been secretly modified by the Soviets to transmit to a nearby listening post all keystrokes typed on those machines. Access to these typewriters to implement the modification was achieved during part of the logistics phase prior to their delivery to the embassy at a point when the typewriters were unsecured.

SOURCE: U.S. House of Representatives, Rep. Henry J. Hyde, Introduction to “Embassy Moscow: Attitudes and Errors,” *Congressional Record*, October 25, 1990, E3489.

—A vendor with an employee loyal to an attacker introduces malicious code as part of a system component for which it is a subcontractor in a critical system.²³

—An attacker intercepts a set of CD-ROMs ordered by the victim and substitutes a different doctored set for actual delivery to the victim. The doctored CD-ROMs contain attack software that the victim installs as he uses the CDs.

—An attacker bribes a shipping clerk to look the other way when the computer is on the loading dock for transport to the victim, opens the box, replaces the video card installed by the vendor with one modified by the attacker, and reseals the box.

- *Compromises of third-party security software.* Security software is intended to protect a computer from outside threats. In many cases, it does so by identifying and blocking specific malicious software or activities based on some kind of “signature” associated with a given malicious action. But a government could induce the vendor of such security software to ignore threats that might be associated with a virus or worm that the government might use to attack an adversary’s system. The government could induce such cooperation in many ways. For example, it could persuade the CEO of the vendor’s company to take such action, prevent the company from selling its products if it failed to take such action, or

²³ For a partial compendium of instances in which vendors have shipped to customers products “pre-infected” with malware (e.g., virus or other security vulnerability or problem), see <http://www.attrition.org/errata/cpo/>.

bribe some low-level programmer to “forget” to include a particular signature in the virus checker’s database.²⁴

- *Compromises in the patch process.* Patching software defects, especially those that fix known vulnerabilities, is an increasingly routine part of system maintenance. Yet patching introduces another opportunity for introducing new vulnerabilities or for sustaining old ones.²⁵ Both automated (e.g., Windows updates) and manual patch processes present opportunities for close-access cyberattacks, though the tools and resources required may be quite different. The patch issued may be corrupted by the cyberattacker; alternatively, the distribution channel itself may be compromised and hostile software installed.

2.2.5.3 Compromise of Operators, Users, and Service Providers

Human beings who operate and use IT systems of interest constitute an important set of vulnerabilities for cyberattack. They can be compromised through recruitment, bribery, blackmail, deception, or extortion. Spies working for the attacker may be unknowingly hired by the victim, and users can be deceived into actions that compromise security. Misuse of authorized access, whether deliberate or accidental, can help an attacker to take advantage of any of the vulnerabilities previously described—and in particular can facilitate close-access cyberattacks.

For example, the operation of a modern nationwide electric power grid involves many networked information systems and human operators of those systems; these operators work with their information systems to

²⁴ Some possible precedent for such actions can be found in the statement of Eric Chien, then chief researcher at the Symantec antivirus research lab, that Symantec would avoid updating its antivirus tools to detect a keystroke logging tool that was used only by the FBI (see John Leyden, “AV Vendors Split over FBI Trojan Snoops,” *The Register*, November 27, 2001, available at http://www.theregister.co.uk/2001/11/27/av_vendors_split_over_fbi/). More discussion of this possibility can be found in Declan McCullagh and Anne Broache, “Will Security Firms Detect Police Spyware?,” *CNET News*, July 17, 2007, available at <http://news.cnet.com/2100-7348-6197020.html>. Other corporate cooperation with government authorities was documented in the Church Committee hearings in 1976. For example, RCA Global and ITT World Communications “provided virtually all their international message traffic to NSA” in the period between August 1945 and May 1975 (see Book III of Supplementary Detailed Staff Reports on Intelligence Activities and the Rights of Americans, 94th Congress, Report 94-755, p. 765). As for a government influencing vendors to compromise security in their products, the canonical example is that for many years, the United States had imposed export controls on information technology vendors selling products with encryption capabilities—allowing more relaxed export controls only on those products capable of weak encryption. Most export controls on strong encryption products were lifted in the late 1990s.

²⁵ Defense Science Board, “Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software,” U.S. Department of Defense, September 2007, p. 25.

keep the system in dynamic balance as loads and generators and transmission facilities come on and off line. A cyberattack might be directed at the systems in a control center that provide situational awareness for the operators—and so operators might not be aware of an emerging problem (perhaps a problem induced by a simultaneous and coordinated physical attack) until it is too late to recover from it.

In many instances involving the compromise of users or operators, the channels for compromise often involve e-mails, instant messages, and files that are sent to the target at the initiative of the attacker, or other sources that are visited at the initiative of the target. Examples of the latter include appealing web pages and certain shareware programs, such as those for sharing music files, or even playing a music CD with rootkit-installation software.

An appealing web page might attract many viewers in a short period of time, and viewers could be compromised simply by viewing the page, while shareware programs might contain viruses or other malware. In an interesting experiment at West Point in 2004, an apparently legitimate e-mail was sent to 500 cadets asking them to click on a link to verify grades. Despite their start-of-semester training (including discussions of viruses, worms, and other malicious code, or malware), over 80 percent of recipients clicked on the link in the message.²⁶

Another example of social engineering in cyberattack involved a red team's use of inexpensive universal serial bus (USB) flash drives to penetrate an organization's security. These drives were scattered in parking lots, smoking areas, and other areas of high traffic. In addition to some innocuous images, each drive was preprogrammed with software that could have collected passwords, logins, and machine-specific information from the user's computer, and then e-mail the findings to the red team. Because many systems support an "auto-run" feature for insertable media (i.e., when the medium is inserted, the system automatically runs a program named "autorun.exe" on the medium) and the feature is often turned on, the red team was notified as soon as the drive was inserted. The result: 75 percent of the USB drives distributed were inserted into a computer.²⁷

A final category of vulnerabilities and access emanates from the IT-based service providers on which many organizations and individuals rely. Both individuals and organizations obtain Internet connectivity from

²⁶ See Aaron J. Ferguson, "Fostering Email Security Awareness: The West Point Carronade," *EDUCAUSE Quarterly* 28(1):54-57, 2005, available at <http://net.educause.edu/ir/library/pdf/EQM0517.pdf>.

²⁷ Steve Stasiukonis, "Social Engineering, the USB Way," *Dark Reading*, June 7, 2006, available at http://www.darkreading.com/document.asp?doc_id=95556&WT.svl=column1_1.

Internet service providers. Many organizations also make use of external firms to arrange employee travel or to manage their IT security or repair needs. These service providers are potential additional security vulnerabilities, and thus might well be targeted in a cyberattack directed at the original organization.

Note: Close-access attacks and social engineering are activities in which national intelligence agencies specialize, and over the years these agencies have presumably developed considerable expertise in carrying out such activities. In practice, it is often cheaper and easier to compromise a person than it is to break through firewalls and decrypt passwords. Indeed, in many situations, human subversion and physical action are the two quickest, cheapest, and most effective methods of attacking a computer system or network.

2.2.6 Propagating a Large-Scale Cyber Offensive Action

In order to take control of a large number of computers, an attacker needs to locate vulnerable computers and somehow install malicious software on those computers. This can be done using direct attacks against exposed services (e.g., scan and attack behavior seen in worms like Slammer and Blaster), or indirectly using social engineering techniques (e.g., e-mail with Trojan horse executables as file attachments, instant messages with hypertext links, web pages containing malicious content, Trojan horse executables in modified “free” software download archives, or removable media devices dropped in parking lots).

2.2.6.1 Direct Propagation

Direct attacks are the fastest means of compromising a large number of hosts. The most common method of direct propagation is either by scanning for vulnerable hosts followed by direct remote attack, or by simply choosing random addresses and attempting to use vulnerabilities regardless of whether there is a host listening on that IP address, or even possessing the vulnerability at all. Malware artifacts (e.g., Agobot/Phatbot²⁸) often look for opportunistic avenues for attack, including the back doors left by other malware that may have previously infected the host. This tactic does not require the use of a new zero-day exploit, as there may be plenty of other commonly known ways to take control of computers. The successful hit rate is lowest by scanning entirely randomly while attacking, although there is an element of surprise using this method because

²⁸ LURHQ, “Phatbot Trojan Analysis,” June 2004, available at <http://www.lurhq.com/phatbot.html>.

there is no opportunity for the target to see the reconnaissance scans. Either way, a direct attack increases the chances of detection through either signature (e.g., IDS/IPS or AV scanning) or anomalous flow detection, possibly triggering a reaction by security operations personnel. Even if an attacker launches a new attack whose signature is not known to the defender, the defender may still be able to use traffic flow analysis to detect the attack by observing and noting anomalous flows.

As mentioned above, worms to date have been quite noisy and in some cases spread so fast that they disrupt the network infrastructure devices themselves. In order to make direct attacks viable to recruit hosts for the kind of attack described here, a more slow and subtle attack (especially one involving a zero-day attack method whose existence is kept secret) over a much longer period of time would be needed.

The methods just described are active in nature, but there are also opportunities for passive direct propagation. For example, hosts infected with the Blaster worm can still be observed actively attempting to propagate, suggesting that attacking through more recently discovered vulnerabilities is likely to be feasible.²⁹ By simply passively monitoring for signs of Blaster-infected hosts scanning randomly across the Internet, one can identify hosts that have a high probability of possessing one or more remotely usable vulnerabilities. Those hosts can then be compromised and taken over, as was done by the Agobot/Phatbot malware in 2003/2004.³⁰ There is also a very good chance that attacks against these hosts, since they were already compromised and are actively scanning the Internet for more victims, would not be noticed by the owners of the computers or networks in which they reside.

2.2.6.2 Indirect Propagation

Indirect methods of cyberattack are often slower, but less easy to detect by either network level IDS/IPS or AV/anti-spam systems. Some indirect methods of cyberattack include:

- Compromising installation archives of freeware or shareware programs, either used generally or known to specifically be used by target organizations. Examples include open source software development efforts, or software linked from sites that aggregate and categorize free-

²⁹ Michael Bailey, Evan Cooke, Farnam Jahanian, David Watson, and Jose Nazario, "The Blaster Worm: Then and Now," *IEEE Security and Privacy* 3(4):26-31, 2005, available at <http://ieeexplore.ieee.org/iel5/8013/32072/01492337.pdf?arnumber=1492337>.

³⁰ Joe Stewart, "Phatbot Trojan Analysis," March 15, 2004, available at <http://www.secureworks.com/research/threats/phatbot/?threat=phatbot>.

ware and shareware programs.³¹ Unwitting users download these altered installation archives and install them without first verifying hashes or cryptographic signatures that attest to their integrity. (In general, use of cryptography for authentication of software is done poorly, or not at all, allowing these kinds of attacks to succeed in many cases.)

- Drive-by download attacks resulting from redirection of clients by corruption of DNS resolver configurations, or man-in-the-middle attacks on DNS requests, for example in free WiFi networks. The attacker who wishes to redirect a software download request, or web page request, must simply answer an unauthenticated DNS request that is easily seen by the attacker in an open WiFi network. The client is then silently redirected to a malicious site, where malicious software is downloaded and installed onto the system.³²

- Cross-site scripting attacks involve redirection of web browsers through embedded content in HTML or Javascript. The redirection is invisible to the user, and can result in portions of a web session being hijacked to install malicious content and/or to capture login credentials that can be used later to compromise the user's account. As an example, some online auctioning sites have a significant problem with their users attacking each other using the site as a platform; the mechanism is that the site permits upload of HTML to its regular vendor's own internal-to-site web pages and in that HTML are hidden various attack mechanisms.

2.2.7 Economics

The economics of cyberattack are quite different from those of kinetic attack.

- *Productivity.* Certain kinds of cyberattacks can be undertaken with much more productivity than can kinetic attacks because the latter depend on human actions while the former depend on computer systems. Automation can increase the amount of damage that can be done per

³¹ An incident related to such compromises was reported in August 2008. The Red Hat Network, distributors for the Linux operating system, detected an intrusion on some of its computer systems. The intruder was able to sign a small number of OpenSSH packages (OpenBSD's SSH (Secure SHell) protocol implementation), and these packages may have been improperly modified. See <https://rhn.redhat.com/errata/RHSA-2008-0855.html>. Some users relying on a Red Hat's digital signature to ensure that they install only authorized software are thus at some potential risk.

³² This is mostly a risk to users who perform normal tasks like reading news on websites from accounts on their computer that have elevated system administrator privileges. For this reason, it is typically recommended that users employ the concept of least-privileges and use accounts with administrator rights only when installing or configuring software, not for general tasks.

attacker, increase the speed at which the damage is done, and decrease the required knowledge and skill level of the operator of the system. (A corollary is that the scale of effects of a cyberattack may be only weakly correlated with effort. As early as 1988, the Morris worm—developed with relatively little effort—demonstrated that a cyberattack could have effects on national and international scales. Indeed, the most difficult part of some cyberattacks may well be to limit their effects to specific computers.) Automation can also simplify operational tasking by providing capabilities such as automated target acquisition, reducing effects to a set of alternatives in a pull-down menu, and turning rules of engagement into operational parameters that tie available actions to targets in the system's menu. The result is a system that is easier to operate than a collection of discrete attack tools and thus requires lower levels of training, knowledge, and specific skills.

- *Capital investment.* A cyberattack architecture functioning on a service-oriented model can be replicated at very low cost if it relies on stolen services. For example, millions of compromised computers assembled in a botnet can be tasked to any C2 control center, allowing a larger number of individual attackers to operate independently. This distributes the operational and management loads, similar to the way a military battalion is composed of companies, cohesive units using similar weapons and tactics but capable of attacking different objective targets at different locations at the same time. Implementing a service-oriented model for a distributed architecture is simply a matter of programming and separation of duties (i.e., acquisition of newly compromised hosts to be controlled, and command and control of subsets of these hosts by individual operational warfare units). The more loosely coupled the functions of command and control versus effects on and from compromised end hosts, the more resistant the overall architecture is to detection and mitigation.

On the other hand, highly specialized cyberweapons (e.g., those designed to attack very specific targets) may well be costly. For example, the development of a particular cyberweapon may require intelligence collection that is difficult and thus expensive to perform. Other cyberweapons may only be useful against adversary targets one or a few times (Section 2.3.10), making their use an expensive proposition.

- *Funding.* Financial assets of an adversary can be used by an attacker. Rather than paying a defense contractor market rates to develop arms and munitions out of its own public coffers, a nation has the ability to steal money from an adversary for use in developing and advancing its cyberattack capabilities. For example, it could develop Version 1.0 of an attack platform on its own and then use the proceeds from fraud perpetrated using Version 1.0 to fund development of a larger and more effective Version 2.0 platform, and so forth. Such an approach could be

particularly appealing to subnational groups, such as terrorist or criminal organizations, or—in the absence of specific legal prohibitions against such actions—to underfunded government agencies.

- *Availability.* The underlying technology for carrying out cyberattacks is widely available, inexpensive, and easy to obtain. Software packages embedding some of the technology for carrying out cyberattacks are available on the Internet, complete with user manuals and point-and-click interfaces. The corollary is that government has no monopoly on cyberweapons or over expertise. Private businesses and private individuals can own or control major cyberweapons with significant capability, but the same tends to be less true of kinetic weapons, citizen-built truck bombs notwithstanding.

2.3 OPERATIONAL CONSIDERATIONS

The previous section addresses the basic technologies of and approaches to cyberattack. This section considers the operational implications of using cyberattack. Both nation-states and hackers must grapple with these implications, but the scope of these implications is of course much broader for the nation-state than for the hacker.

2.3.1 The Effects of Cyberattack

Although the ultimate objective of using any kind of weapon is to deny the adversary the use of some capability, it is helpful to separate the effects of using a weapon into its direct and its indirect effects (if any). The direct effects of using a weapon are experienced by its immediate target. For example, the user of a kinetic weapon seeks to harm, damage, destroy, or disable a physical entity. The indirect effects of using that weapon are associated with the follow-on consequences of harming, damaging, destroying, or disabling a physical entity, which may include harming, destroying, or disabling *other* physical entities—a runway may be damaged (the direct effect) so that aircraft cannot land or take off (the indirect effect). This distinction between direct and indirect effects is particularly important in a cyberattack context.

2.3.1.1 Direct Effects³³

By definition, cyberattacks are directed against computers or networks. The range of possible direct targets for a cyberattack is quite broad and includes (but is not limited to) the following:

³³Much of the discussion in this section is based on National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington D.C., 2007.

- *Computer chips embedded in other devices*, such as weapons systems, communications devices, generators, medical equipment, automobiles, elevators, and so on. In general, these microprocessors provide some kind of real-time capability (e.g., a supervisory control and data acquisition system will control the operation of a generator or a floodgate, a chip in an automobile will control the flow of fuel, a chip in an ATM will dispense money).

- *The computing systems controlling elements of the nation's critical infrastructure*, for example, the electric power grid, the air traffic control system, the transportation infrastructure, the financial system, water purification and delivery, or telephony. For example, cyberattacks against the systems and networks that control and manage elements of a nation's transportation infrastructure could introduce chaos and disruption on a large scale that could drastically reduce the capability for transporting people and/or freight (including food and fuel).

- *Dedicated computing devices* (e.g., desktop or mainframe computers). Such devices might well not be just any desktop computer (e.g., any computer used in offices around the country) but rather the desktop computers in particular sensitive offices, or in critical operational software used in corporate or government computer centers (e.g., a major bank or the unclassified systems of an adversary nation's ministry of defense). Dedicated computer systems might also include the routers that control and direct traffic on the Internet or on any other network.

Cyberattacks generally target one of several attributes of these components or devices—they seek to cause a loss of integrity, a loss of authenticity, or a loss of availability (which includes theft of services):

- *Integrity*. A secure system produces the same results or information whether or not the system has been attacked. An attack on integrity seeks to alter information (a computer program, data, or both) so that under some circumstances of operation, the computer system does not provide the accurate results or information that one would normally expect even though the system may continue to operate. A computer whose integrity has been compromised might be directed to destroy itself, which it could do if it were instructed to turn off its cooling fan. A loss of integrity also includes suborning a computer for use in a botnet, discussed further in Section 2.2.5.1.1.

- *Authenticity*. An authentic message is one that is known to have originated from the party claiming to have originated it. An attack on authenticity is one in which the source of a given piece of information is obscured or forged. A message whose authenticity has been compromised will fool a recipient into thinking it was properly sent by the asserted originator.

- *Availability.* A secure system is available for normal use by its rightful owner even in the face of an attack. An attack on availability may mean that e-mail sent by the targeted user does not go through, or the target user's computer simply freezes, or the response time for that computer becomes intolerably long (possibly leading to catastrophe if a physical process is being controlled by the system). Some analysts also discuss theft of services—an adversary may assume control of a computer to do his bidding. In such situations, the availability of the system has been *increased*, but for the wrong party (namely, the adversary).

These attributes may be targeted separately or together. For example, a given cyberattack may support the compromise of both integrity and availability, though not necessarily at the same time. In addition, the victim may not even be aware of compromises when they happen—a victim may not know that an attacker has altered a crucial database, or that he or she does not have access to a particular seldom-used emergency system.

In some situations, integrity is the key target, as it might well be for a tactical network. A commander who doubts the trustworthiness of the network used to transmit and receive information will have many opportunities for second-guessing himself, and the network may become unreliable for tactical purposes. In other situations, authenticity is the key target—a cyberattack may take the form of a forged message purportedly from a unit's commanders to move from one location to another. And in still other situations, availability is the target—a cyberattack may be intended to turn off the sensors of a key observation asset for the few minutes that it takes for kinetic assets (e.g., airplanes) to fly past it.

The direct effects of some cyberattacks may be easily reversible. (Reversibility means that the target of the attack is restored to the operating condition that existed prior to the attack.) For example, turning off a denial-of-service attack provides instant reversibility with no effort on the part of the attacked computer or its operators. If backups are available, an attack on the integrity of the operating system may take just a few minutes of reloading the operating system. Many effects of kinetic attacks are not as easy to reverse.³⁴

A corollary to this point is that achieving enduring effects from a cyberattack may require repeated cyberstrikes, much as repeated bombing of an airstrip might be necessary to keep it inactive. If so, keeping a

³⁴ For example, the time scales involved may be very different. Restoring the capability of an attacked computer that controls a power distribution system is likely to be less costly or time-consuming compared to rebuilding a power plant damaged by kinetic weapons. (A cyberattack on a computer controlling a power distribution system may even be intended to give the attacker physical control over the system but not to damage it, enabling him to control production and distribution as though he were authorized to do so.)

targeted system down is likely to be much more difficult than bringing it down in the first place, not least because the administrators of the victimized system will be guided by the nature of the first attack to close off possible attack paths. Thus, the attacker may have to find different ways to attack if the goal is to create continued effects. That is, depending on the nature of his goals, the attacker must have operational plans that anticipate possible defense actions and identify appropriate responses should those defense actions occur.

2.3.1.2 Indirect (and Unintended) Effects

Although the direct effects of a cyberattack relate to computers, networks, or the information processed or transmitted therein, cyberattacks are often launched in order to obtain some other, indirect effect—and in no sense should this indirect effect be regarded as secondary or unimportant. The adversary air defense radar controlled by a computer is of greater interest to the U.S. commander in the field than is the computer itself. The adversary's generator controlled by a computer is of greater interest to the U.S. National Command Authority than is that computer itself.³⁵ In such cases, the indirect effect is more important than the first-order direct effect.

Computers are also integral parts of command and control networks. In this case, the indirect effect sought by compromising such a computer is to prevent or delay the transmission of important messages, or to alter the contents of such messages.

Indirect effects—which are often the primary goal of a cyberattack—are generally not reversible. A cyberattack may disrupt a computer that controls a generator. The attack on the computer may be reversible (leaving the computer as good as new). But the follow-on effect—the generator overheating and destroying itself—is not reversible.

Cyberattacks are particularly well suited for attacks on the psychology of adversary decision makers who rely on the affected computers, and in this case such effects can be regarded as indirect effects. For example, a single database that is found to be deliberately corrupted, even when controls are in place to prevent such corruption, may call into question the integrity of all of the databases in a system. It is true that all production databases have some errors in them, and indeed savvy users ought to

³⁵ For example, a test staged by researchers at the Department of Energy's Idaho National Laboratories used a cyberattack to cause a generator to self-destruct. Known as Aurora, the cyberattack was used to change the operating cycle of the generator, sending it out of control. See CNN, "Staged Cyber Attack Reveals Vulnerability in Power Grid," September 26, 2007, available at <http://www.cnn.com/2007/US/09/26/power.at.risk/index.html>.

adjust for the possibility that their data may be incorrect. But in practice, they often do not. Being made conscious of the fact that a database may have been compromised has definite psychological effects on a user. Thus, the victim may face a choice of working with data that may—or may not—have been corrupted and suffering all of the confidence-eroding consequences of working in such an environment, or expending enormous amounts of effort to ensure that other databases have not been corrupted or compromised.³⁶ A second example might be the clandestine alteration of critical data that causes decision makers to make poor or unfavorable decisions.

The unintended consequences of a cyberattack are almost always indirect effects. For example, a cyberattack may be intended to shut down the computer regulating electric power generation for a Zendian air defense facility. The direct effect of the cyberattack could be the disabling of the computer. The intended indirect effect is that the air defense facility loses power and stops operating. However, if—unknown to the attacked—a Zendian hospital is also connected to the same generation facility, the hospital's loss of power and ensuing patient deaths are also indirect effects, and also an unintended consequence, of that cyberattack.

2.3.2 Possible Objectives of Cyberattack

Whether a cyberattack is conducted remotely or through close access, what might it seek to accomplish? Some possible objectives include the following, in which an attacker might seek to:

- *Destroy a network or a system connected to it.* Destruction of a network or of connected systems may be difficult if “destruction” means the physical destruction of the relevant hardware, but is much easier if “destruction” simply means destroying the data stored within and/or eliminating the application or operating systems programs that run on that hardware. For example, an attacker might seek to delete and erase permanently all data files or to reformat and wipe clean all hard disks that it can find. Moreover, destruction of a network also has negative consequences for anything connected to it—power-generation facilities

³⁶ For example, in 1982, the United States was allegedly able to “spike” technology that was subsequently stolen by the Soviet Union. Above and beyond the immediate effects of its catastrophic failure in a Soviet pipeline, Thomas Reed writes that “in time the Soviets came to understand that they had been stealing bogus technology, but now what were they to do? By implication, every cell of the Soviet leviathan might be infected. They had no way of knowing which equipment was sound, which was bogus. All was suspect, which was the intended endgame for the entire operation.” See Thomas C. Reed, *At the Abyss: An Insider's History of the Cold War*, Ballantine Books, New York, 2004.

controlled by a network are likely to be adversely affected by a disabled network, for example.

- *Be an active member of a network and generate bogus traffic.* For example, an attacker might wish to masquerade as the adversary's national command authority or as another senior official (or agency) and issue phony orders or pass faked intelligence information. Such impersonation (even under a made-up identity) might well be successful in a large organization in which people routinely communicate with others that they do not know personally. Alternatively, the attacker might pretend to be a non-existent agency within the adversary's government and generate traffic to the outside world that looks authentic. An impersonation objective can be achieved by an attacker taking over the operation of a trusted machine belonging to the agency or entity of interest (e.g., the National Command Authority) or by obtaining the relevant keys that underlie their authentication and encryption mechanisms and setting up a new node on the network that appears to be legitimate because it exhibits knowledge of those keys.

- *Clandestinely alter data in a database stored on the network.* For example, the logistics deployment plan for an adversary's armed forces may be driven by a set of database entries that describe the appropriate arrival sequence of various items (food, fuel, vehicles, and so on). A planner relying on a corrupted database may well find that deployed forces have too much of certain items and not enough of others. The planner's confidence in the integrity of the database may also be affected, as discussed in Section 2.3.1.2.

- *Degrade or deny service on a network.* An attacker might try to degrade the quality of service available to network users by flooding communications channels with large amounts of bogus traffic—spam attacks can render e-mail ineffective as a communications medium, for example. Denial-of-service attacks might be directed at key financial institutions, for example, and greatly degrade their ability to handle consumer financial transactions. A denial-of-service attack on the wireless network (e.g., a jamming attack) used to control a factory's operations might well shut it down. Taking over a telecommunications exchange might give an attacker the ability to overwhelm an adversary's ministry of defense with bogus phone calls and make it impossible for its employees to use its telephones to do any work. A denial-of-service attack might be used to prevent an adversary from using a communications system, and thereby force him to use a less secure method for communications against which a cyber-exploitation could be successful.

- *Assume control of a network and/or modulate connectivity, privileges, or service.* An attacker might assume control of an Internet service provider in an adversary nation, and decide who would get what services and con-

nectivity. For example, it might intentionally (and clandestinely) degrade bandwidth to key users served by that ISP, so that transmission of large files (e.g., images or video) would take much longer than expected. If the ISP was used by the Zendian Ministry of Defense to pass targeting information for prompt action, delays in transmission might cause Zendian forces to miss important deadlines.

Finally, cyberattacks can be carried out in conjunction with kinetic attacks, and indeed the effect of a cyberattack may be maximized if used in such a manner. For example, a cyberattack alone might be used to cause confusion. But the ability to cause confusion in the midst of a kinetic attack might well have greater operational significance for the attacker.

2.3.3 Target Identification

As with any other weapon, a cyberattack must be directed at specific computers and networks. Even if a nation-state has been identified as being subject to cyberattack, how can the specific computers or networks of interest be identified in a useful manner? (Note also that target identification is often related to attribution, but does not necessarily follow—a computer posing a threat may well be regarded as a target, even if the party controlling it is not known.)

In some instances, the target identification process is a manual, intelligence-based effort. From a high-level description of the targets of interest (e.g., the vice president's laptop, the SCADA systems controlling the electric generation facility that powers the air defense radar complex 10 miles north of the Zendian capital, the transaction processing systems of the Zendian national bank), a route to those specific targets must be found. For example, a target system with an Internet presence may have an IP address. Knowledge of the system's IP address provides an initial starting point for attempting to gain entry to the appropriate component of the target system.

Sometimes a computer is connected to the Internet indirectly. For example, although it is common for SCADA systems to be putatively "air gapped" from the Internet, utility companies often connect SCADA systems to networks intended for administrative use so that their business units can have direct real-time access to the data provided by SCADA systems to improve efficiency.³⁷ Compromising a user of the administrative

³⁷ For example, in January 2003, the Slammer worm downed one utility's critical SCADA network after moving from a corporate network, through a remote computer to a VPN connection to the control center LAN. See North American Electric Reliability Council, "SQL Slammer Worm Lessons Learned for Consideration by the Electricity Sector," June 20, 2003, available at http://www.esisac.com/publicdocs/SQL_Slammer_2003.pdf.

network may enable an attacker to gain access to these SCADA systems, and thus intelligence collection efforts might focus on such users.

Target identification information can come from a mix of sources, including open source collection, automated probes that yield network topology, and manual exploration of possible targets. Manual target identification is slow, but is arguably more accurate than automated target identification.

Automated target selection is based on various methods of mapping and filtering IP addresses and/or DNS names, for example through programmed pattern matching, network mapping, or querying databases (either public ones, or ones accessible through close-access attacks). The scope of automated attack identification can be limited by the use of network address filtering. Say, for example, an attacker wishes to target a specific military base, or to affect only hosts within a specific corporate network. Various public records exist, such as WHOIS network registration information, BGP network routing information, DNS zone files, and so on, that map Internet-accessible domain names to IP network address blocks.

Automated target selection within an internal network is more complicated. An internal network may have one gateway to the Internet, but within the perimeter of the internal network may be any arrangement of internal addresses. Once an attacker gains access to a host inside the network, the internal DNS zone tables can be accessed and decoded to identify appropriate targets. This will not always be possible, but in many cases even internal network ranges can be determined with minimal effort by the attacker. It is also possible to perform simple tests, such as attempting to access controlled websites to test the ability to make outbound (i.e., through the firewall) connections³⁸ and thus to determine network membership through the resulting internal/external address mappings. If the attacker has sufficient lead time, a “low and slow” network probe can—without arousing suspicion—generally yield connectivity information that is adequate for many attack purposes.

A cyberattacker may also be interested in selecting targets that fall under a number of administrative jurisdictions. As a rule of thumb, organizations under different jurisdictions are less willing to share information among themselves than if only a single jurisdiction is affected—and thus a coordinated response to a cyberattack may be less effective than it might otherwise be. Furthermore, different administrative jurisdictions are likely to enforce a variety of security precautions, suggesting that some jurisdictions would be less resistant to an attack than others.

³⁸ An illustration is the use of a query to the domain name system as a covert channel.

2.3.4 Intelligence Requirements and Preparation

Attacks on the confidentiality, integrity, authenticity, and availability attributes require taking advantage of some vulnerability in the targeted system. However, an attacker seeking to exploit a given vulnerability must know—in advance of the attack—whether the targeted system is in fact vulnerable to any particular attack method of choice. Indeed, the success of a cyberattack (to include both achieving the intended goal and minimizing collateral damage) often depends on many details about the actual configuration of the targeted computers and networks.

As a general rule, a scarcity of intelligence information regarding possible targets means that any cyberattack launched against them can only be a “broad-spectrum” and relatively indiscriminate or blunt attack. (Such an attack might be analogous to the Allied strategic bombing attacks of World War II that targeted national infrastructure on the grounds that the infrastructure supported the war effort of the Axis.) Substantial amounts of intelligence information about targets (and paths to those targets) are required if the attack is intended as a very precise one directed at a particular system and/or if the attack is to be a close-access attack.³⁹ Conversely, a lack of such information will result in large uncertainties about the direct and indirect effects of a cyberattack, and make it difficult for commanders to make good estimates of likely collateral damage.

Information collection for cyberattack planning differs from traditional collection for kinetic operations in that it may require greater lead time and may have expanded collection, production, and dissemination requirements, as specific sources and methods may need to be positioned and employed over time to collect the necessary information and conduct necessary analyses.⁴⁰ As illustrations (not intended to be exhaustive), intelligence information may be required on:

- The target’s platform, such as the specific processor model;
- The platform’s operating system, down to the level of the specific version and even the history of security patches applied to the operating system;
- The IP addresses of Internet-connected computers;
- The specific versions of systems administrator tools used;

³⁹ To some extent, similar considerations apply to the intelligence required to support precise kinetic weaponry. If a kinetic weapon is intended to be used, and capable of being used, in a precise manner, more information about the target and its environment will be necessary than if the weapon is blunt and imprecise in its effects.

⁴⁰ Joint Chiefs of Staff, *Information Operations*, Joint Publication 3-13, U.S. Department of Defense, Washington, D.C., February 2006, available at www.dtic.mil/doctrine/jel/new_pubs/jp3_13.pdf.

- The security configuration of the operating system, e.g., whether or not certain services are turned on or off, or what antivirus programs are running;
- The physical configuration of the hardware involved, e.g., what peripherals or computers are physically attached;
- The specific operators of the systems in question, and others who may have physical access to the rooms in which the systems are kept;
- The name of the shipping company that delivers computer components to the facility;
- The telephone numbers of the help desk for the system in question; and so on.

That is, the list of information items possibly relevant to cyberattacks in general is quite long indeed. Although not all such information will be necessary for any given attack, some of it will surely be needed depending on the precise nature of the cyberattack required.

In some cases, such information may be available from public sources, and these sources are often voluminous with a wealth of useful information. In other cases, the required information may not be classified but may be available only through non-official sources, such as in the non-shredded office trash of the installation in question. In still other cases, the relevant information may be available through the traditional techniques of human agents infiltrating an organization or interviewing people familiar with the organization. Finally, automated means may be used to obtain necessary intelligence information—an example is the use of automated probes that seek to determine if a system has ports that are open, accessible, and available for use.

Intelligence preparation for a cyberattack is often a staged process. For example, stolen login credentials can be used to gain access to compromised accounts, followed by escalation of privileges to (a) locate and exfiltrate files or (b) gain complete control over the host, allowing further keystroke logging or password extraction to compromise not only other accounts on the same system, but also accounts on other hosts. This in turn extends the attacker's reach (as seen in the Stakkato case described in Appendix C) and enables the attacker to gather more information that might support an attack. Maximizing the ability to take advantage of these stolen credentials becomes a matter of database entry and processing to automate the tasks.

A cyberattacker will also benefit from knowledge about the adversary's ability to respond in a coordinated manner to a widespread attack. At the low end of this continuum (Table 2.2), the adversary is only minimally capable of responding to attacks even on isolated or single systems, and has no capability at all to take coordinated action against attacks on

TABLE 2.2 Levels of Intrusion Response

Level	Victim Posture	Characteristic Actions
0	Unaware	None: Passive reliance on inherent software capabilities
1	Involved	Uses and maintains antivirus software and personal firewalls
2	Interactive	Modifies software and hardware in response to detected threats
3	Cooperative	Implements joint traceback with other affected parties
4	Non-cooperative (active response)	Implements invasive tracebacks, cease-and-desist measures, and other actions up to retaliatory counterstrikes

SOURCE: David Dittrich, *On the Development of Computer Network Attack Capabilities*, work performed for the National Research Council under agreement D-235-DEPS-2007-001, February 3, 2008.

multiple systems. At the high end of this continuum, the adversary can integrate information relating to attacks on all of the systems under its jurisdiction, develop a relatively high degree of situational awareness, and respond in an active and coordinated manner.⁴¹

Ultimately, the operational commander must make an assessment about whether the information available is adequate to support the execution of a cyberattack. Such assessments are necessarily informed by judgments about risk—which decreases as more information is available. Unfortunately, there is no objective, quantitative way to measure the adequacy of the information available, and also no way to quantitatively ascertain the increase in risk as the result of less information being available (the discussion in Section 2.3.6 elaborates on sources of uncertainty). In practice, the best way to adapt to a lack of detailed information may be to ensure the availability of skilled and adaptive personnel who can modify an attack as it unfolds in response to unanticipated conditions.

Lastly, the fact that considerable intelligence information may be required to conduct a specific targeted attack points to a possible defen-

⁴¹ Even when the capacity and resources exist to be able to operate at a high response level, there are many reasons why system owners may not respond in a cooperative manner to a widespread computer attack. They may not be capable of immediately responding, may lack adequate resources, may be unable to physically attend to the compromised host, or may even speak a different language than the person reporting the incident. There may be active complicity within the organization, or a willful disregard of reports that allow the attacker to continue unabated.

sive strategy—if one knows that a cyberattack is imminent, a defender may take steps to invalidate the intelligence information that the attacker may have collected. Such steps may be internally originated (e.g., changing one’s own system configuration and defensive posture) or externally originated (e.g., downloading a security update provided by a vendor). If these steps are successful (and it may well be possible to change defensive postures rapidly), such action may force the attacker to postpone or abandon his attack or to conduct an attack that is much less precise and focused and/or much less certain in outcome than it would otherwise have been. (These points are far less relevant if the attacker is interested in a “general” attack against broad swaths of an adversary’s computers and networks—in such an attack, the targets of interest are, by definition, the most weakly defended ones.)

2.3.5 Effects Prediction and Damage Assessment

In the kinetic world, weapons (or, more precisely, munitions) are aimed against targets. Predicting the effect of a weapon on a given target is obviously important to operational planners, who must decide the most appropriate weapons-to-target matching. In general, characteristics of the weapon, such as explosive yield, fusing, likely miss distances (more precisely, Circular Error Probable—the distance from the target within which the weapon has a 50 percent chance of striking), and so on are matched against characteristics of the target (such as target hardness, size, and shape), and its surrounding environment (e.g., terrain and weather).

Damage assessment for physical targets is conceptually straightforward—one can generally know the results of a strike by visual reconnaissance, although a task that is straightforward in principle may be complicated by on-the-ground details or adversary deception. For example, the weather may make it impossible to obtain visual imagery of the target site, or the adversary may be able to take advantage of the delay between weapons impact and damage assessment to create a false impression of the damage caused.

There are similar needs for understanding the effect of cyberweapons and assessing damage caused by cyberweapons. But munitions effects and damage assessment are complex and difficult challenges, because the effectiveness of cyberweapons is a strong function of the intelligence available.

Munitions effects in the kinetic world can often be calculated on the basis of computational models that are based on physics-based algorithms. That is, the fundamental physics of explosives technology and of most targets is well known, and so kinetic effects on a given target can be calculated with acceptable confidence. Thus, many of the uncertainties in

kinetic targeting can be theoretically calculated and empirically validated (e.g., at weapons effects test ranges), and the remaining uncertainties relate to matters such as target selection and collocation of other entities with the intended target.

But there is no comparable formalism for understanding the effects of cyberweapons. The smallest change in the configuration and interconnections of an IT system can result in completely different system behavior, and the direct effects of a cyberattack on a given system may be driven by the behavior and actions of the human system operator and the specific nature of that system as well as the intrinsic characteristics of the cyberweapon involved. Furthermore, these relatively small and/or obscure and/or hidden characteristics are often important in cyber targeting, and information about these things is difficult to obtain through remote intelligence collection methods such as photo reconnaissance, which means that substantial amounts of relevant information may not be available to the attacker.

An example of an error causing unexpected behavior in a cyberattack is the Sapphire/Slammer worm of January 2003. Although the Sapphire worm was the fastest computer worm in history (infecting more than 90 percent of vulnerable hosts within 10 minutes), a defective random number generator significantly reduced its rate of spread.⁴² (The worm targeted IP addresses chosen at random, and the random number generator produced numbers that were improperly restricted in range.) In a military attack context, a cyberattack that manifested its effects more slowly than anticipated might be problematic.

An additional complication to the prediction problem is the possibility of cascading effects that go further than expected. For example, in analyzing the possible effects of a cyberattack, there may be no good analog to the notion of a lethal radius within which any target will be destroyed. When computer systems are interconnected, damage to a computer at the NATO Defense College in Italy can propagate to a computer at the U.S Air Force Rome Laboratory in New York—and whether or not such a propagation occurs depends on a detail as small as the setting on a single switch, or the precise properties of every device connected at each end of the link, or the software characteristics of the link itself.

Engineers often turn to test ranges to better understand weapons effects, especially in those instances in which a good theoretical understanding is not available. A weapons test range provides a venue for testing weapons empirically—sending them against real or simulated targets and observing and measuring their effects. Such information, suitably

⁴² David Moore, “The Spread of the Sapphire/Slammer Worm,” undated publication, available at <http://www.caida.org/publications/papers/2003/sapphire/sapphire.html>.

refined, is then made available to users to assist them in the weapons selection process.

A certain kind of cyberweapon may need to be tested against different versions of operating systems (and indeed, even against different builds of the same operating system), different configurations of the same operating system, and even against different operators of that system. To test for cascading effects, multiple computers must be interconnected. Thus, realistic test ranges for cyberweapons are inevitably complex. It is also quite difficult to configure a cyber test range so that a simulation will provide high confidence that a given cyberattack will be successful.

Some analysts take from these comments that the effects of a cyberattack are impossible to predict. As a blanket statement, this claim is far overstated. It is true that the launch of a worm or virus may go on to infect millions of susceptible computers, and some of the infected machines might happen to control an electric power grid or a hospital information system. The media often report such events as if they were a surprise—and indeed it may well have been a surprise that these particular machines were infected. Nevertheless, after-the-fact analysis of such cyberattacks sometimes leads to the conclusion that the party launching the attack could have predicted the number of susceptible machines fairly accurately.⁴³

Indeed, more customized cyberattacks are quite possible, depending on the *goal* of the attacker. A software agent introduced into a target machine could, in principle, search its environment and remain resident only if that search found certain characteristics (e.g., if the machine had more than 10 files containing the phrases “nuclear weapon” and “Washington D.C.” and had an IP address in a particular range, which might translate into the nation in which the machine was located).

Nevertheless, high degrees of customization may require large amounts of information on machine-identifiable characteristics of appropriate targets. Such information may be obtained through traditional intelligence collection methods. In some cases, a scouting agent undertakes the initial penetration, explores the targeted machine or network to obtain the necessary information, and then retrieves the appropriate exploit from its controller to carry out the next step in the attack.

To illustrate, the precise geographical location of a computer is often not available to a software agent running on it, and may indeed be impos-

⁴³ These comments presume that the attack software is written correctly as the attacker intended—mistakes in the worm or virus may indeed lead to unintended effects. A classic example of a worm written with unintended consequences is the Morris worm. See Brendan P. Kehoe, “Zen and the Art of the Internet,” available at http://www.cs.indiana.edu/docproject/zen/zen-1.0_10.html#SEC91.

sible for the agent to discover. On the other hand, its topological relationship to other systems may be the variable of significance, which may or may not be as good an indicator of function or importance, and topological relationships are likely to be discoverable by an agent.

An issue for uncustomized cyberattacks is “blowback,” which refers to a bad consequence returning to the instigator of a particular action. In the cyberattack context, blowback may refer to direct damage caused to one’s own computers and networks as the result of a cyberattack that one has launched. For example, if the United States launched a cyberattack against an adversary using a rapidly multiplying but uncustomized worm over the Internet, the worm might return to adversely affect U.S. computers and networks. It might also refer to indirect damage—a large-scale U.S. cyberattack against a major trading partner’s economic infrastructure might have effects that could harm the U.S. economy as well.

Another class of weapons effects might be termed strategic. Tactical effects of using a weapon manifest themselves immediately and generally involve destruction, disabling, or damage of a target—tactical attacks seek an immediate effect on an adversary and its military forces. By contrast, strategic effects are less tangible and emerge over longer periods of time—strategic attacks are directed at adversary targets with the intent or purpose of reducing an adversary’s warmaking capacity and/or will to make war against the United States or its allies, and are intended to have a long-range rather than an immediate effect on an adversary. Strategic targets include but are not limited to key manufacturing systems, sources of raw material, critical material, stockpiles, power systems, transportation systems, communication facilities, and other such systems.⁴⁴

Most importantly, strategic effects are often less predictable than tactical effects. For instance, recall the history of the German bombing of London in World War II. Originally believed by the Germans to be (among other things) a method of reducing civilian support for the British government, it proved to have the opposite effect.⁴⁵ As a new modality of offensive action, the strategic impact of cyberattack on a population would be even harder to predict in the absence of empirical evidence.

As for assessing damage caused by a cyberattack, note first that the damage due to a cyberattack is usually invisible to the human eye. To ascertain the effects of a computer network attack over the Internet, an

⁴⁴ See DOD definitions for “strategic operations” and “strategic air warfare,” in Joint Chiefs of Staff, *Dictionary of Military and Associated Terms*, Joint Publication 1-02, Department of Defense, Washington, D.C., April 12, 2001 (as amended through October 17, 2008), available at http://www.dtic.mil/doctrine/jel/new_pubs/jp1_02.pdf.

⁴⁵ See, for example, L. Morgan Banks and Larry James, “Warfare, Terrorism, and Psychology,” pp. 216-222 in *Psychology of Terrorism*, Bruce Bongar et al., eds., Oxford University Press, 2007.

attacker might be able to use Internet tools such as *ping* and *traceroute*. These tools are commonly available, and they test the routing from the sending machine to the target machine. *Ping* is intended to measure the round-trip time for a packet to travel between the sending machine and the target machine—and if the target machine is down, *ping* will return an error message saying it could not reach the target machine. *Traceroute* reports on the specific path that a given packet took from the sending machine to the target machine—and if the target machine is down, *traceroute* returns a similar error message. Thus, the receipt of such an error message by the attacker may indicate that an attack on the target has been successful. But it also may mean that the operators of the target machine have turned off the features that respond to *ping* and *traceroute*, so as to lower their vulnerability to criminal hackers or to mislead the damage assessor about the effectiveness of an attack.

More generally, a cyberattack is—by definition—intended to impair the operation of a targeted computer or network. But from a distance, it can be very difficult to distinguish between the successful outcome of a cyberattack and a faked outcome. For example, an attack may be intended to disrupt the operation of a specific computer. But the attacker is faced with distinguishing between two very different scenarios. The first is that the attack was successful and thus that the targeted computer was disabled; the second is that the attack was unsuccessful and also was discovered, and that the adversary has turned off the computer deliberately—and can turn it on again at a moment's notice.

How might this problem be avoided? Where *ping* and *traceroute* as tools for damage assessment depend on the association of a damaged machine with a successful cyberattack, an alternative approach might call for the use of in-place sensors that can report on the effects of a cyberattack. Prior to a cyberattack intended to damage a target machine, the attacker plants sensors of its own on the target machine. These sensors respond to inquiries from the attacker and are programmed to report back to the attacker periodically. These sensors could also be implanted at the same time as attack capabilities are installed on the target machine. Such sensors could report on the outcomes of certain kinds of cyberattacks, and thus in some situations could reduce the uncertainty of damage assessment.

It may also be possible to use non-cyber means for damage assessment of a cyberattack. For example, if a cyberattack is intended to cause a large-scale power outage in a city, its success or failure may be observable by an individual in that city reporting back to the attackers via satellite phone or by an indigenous news network reporting on events within the country. But if the intent of the cyberattack is to turn off the power to a specific radar installation in the nation's air defense network

at a specific time, it will be difficult to distinguish between a successful attack and a smart and wily defender who has detected the attack and shut the power down himself but who is prepared to turn it back on at a moment's notice.

The bottom line on damage assessment is that the state of the art in damage assessment techniques for cyberattacks is still primitive in comparison to similar techniques for kinetic attacks. Cyberattackers must therefore account for larger amounts of uncertainty in their operational planning than their physical-world counterparts—and thus may be inhibited from relying solely or primarily on cyberattack for important missions.

2.3.6 Complexity, Information Requirements, and Uncertainty

From an analytical perspective, it is helpful to separate the uncertainty of effects resulting from cyberattack into several components:

- *Uncertainties that result from any military operation using any kind of weapon.* All military operations are uncertain in outcome to some extent, and all run some risk of incurring collateral damage or having unintended consequences. The availability of intelligence information that is more accurate and more complete reduces the uncertainty inherent in an operation, and it is likely that the necessary intelligence for cyber targets will be less available than for most kinetic targets.

- *Uncertainties that result from the lack of experience with a new weapon.* Additional uncertainties arise when new weapons are used because their operational effects are not well known or well understood. For example, the actual death tolls associated with the Hiroshima and Nagasaki bombings far exceeded the predicted tolls because only blast effects were taken into consideration. (Scientists understood that nuclear weapons had effects other than blast, but they did not know how to estimate their magnitude.) The same has been true for most of the history of U.S. planning for nuclear strikes.⁴⁶

- *Uncertainties that result from unanticipated interactions between civilian and military computing and communications systems.* Because much of the IT infrastructure is shared for military and civilian purposes, disruptions to a military computer system or network may propagate to a civilian system or network. (In some cases, an adversary may *deliberately* intermingle military and civilian assets in order to dissuade the attacker from attacking

⁴⁶ Lynn Eden, *Whole World on Fire: Organizations, Knowledge, and Nuclear Weapons Devastation*, Cornell University Press, Ithaca, N.Y., 2004.

and thereby causing additional collateral damage.⁴⁷) But without detailed knowledge of the interconnections between military and civilian systems, cascading effects may well occur.

As a rule, planning for cyberattack can involve a much larger range of choices and options than planning for most traditional military operations. For example, cyberattack planners must account for a wide range of time and space dimensions. The relevant time scales can range from tenths of a second (a cyberattack may interfere with the timing of a real-time process control system) to years (a cyberattack may seek to implant “sleeper” capabilities in an adversary network that might be activated many years hence). And the systems targeted may be dispersed around the globe or concentrated in a facility next door. All of these factors increase the complexity of planning a cyberattack.

One of the most difficult-to-handle aspects of a cyberattack is that in contrast to a kinetic attack that is almost always intended to destroy a physical target, the desired effects of a cyberattack are almost always indirect, which means that what are normally secondary effects are in fact of central importance. In general, the planner must develop chains of causality—do *X*, and *Y* happens, which causes *Z* to happen, which in turn causes *A* to happen. Also, many of the intervening events between initial cause and ultimate effect are human reactions (e.g., in response to an attack that does *X*, the network administrator will likely respond in way *Y*, which means that *Z*—which may be preplanned—must take response *Y* into account). Moreover, the links in the causal chain may not all be of similar character—they may involve computer actions and results, or human perceptions and decisions, all of which combine into some outcome.

Understanding secondary and tertiary effects often requires highly

⁴⁷ The same is true in reverse—a cyberattack on the civilian side may well result in negative effects on military computers. This point was illustrated by the “I Love You” virus (also referred to as the “Love Bug”), released in May 2000. Press releases and articles from the Department of Defense indicate that some unclassified DOD systems, and even a few classified systems, were infected with this virus. See Jim Garamone, “Love Bug Bites DoD, Others,” American Forces Press Service, May 4, 2000, available at <http://www.defenselink.mil/news/newsarticle.aspx?id=45220>; “Statement by Assistant Secretary of Defense (Public Affairs) Ken Bacon,” U.S. Department of Defense, May, 2000, available at http://findarticles.com/p/articles/mi_pden/is_200005/ai_2892634075. Testimony to a congressional subcommittee from the Government Accounting Office shortly after the virus struck noted the impacts to DOD and many other federal agencies, in addition to the impacts on the private sector. See General Accounting Office, “‘ILOVEYOU’ Computer Virus Highlights Need for Improved Alert and Coordination Capabilities—Statement of Jack L. Brock, Jr.,” Testimony Before the Subcommittee on Financial Institutions, Committee on Banking, Housing and Urban Affairs, U.S. Senate, GAO/T-AIMD-00-181, May 18, 2000.

specialized knowledge. For example, an attack on an electric power grid will require detailed knowledge about the power generation plant of interest—model numbers, engineering diagrams, and so on. Thus, planning a cyberattack may also entail enormous amounts of intellectual coordination among many different individuals (likely scattered through many organizations).

The result is often a complex execution plan, and complex execution plans have many ways to go wrong—the longer a causal chain, the more uncertain its ultimate outcomes and conclusions. This is not simply a matter of unintended consequences of a cyberattack, though that is certainly a concern. The point also relates to the implications of incomplete or overlooked intelligence. For example, it may be that a cyberattack is entirely successful in disabling the computer controlling an air defense radar, but also, as it turns out, that there is a backup computer in place that was not mentioned in the intelligence reports used to plan the attack. Or a connection between two systems that is usually in place is disconnected on the day of the attack because of a maintenance schedule that was changed last week, and thus was unknown to the attack planners—resulting in the inability of the attacker to destroy the backup computer.

One way of coping with uncertainty in this context is to obtain feedback on intermediate results achieved through monitoring the execution of the initial plan and then applying “mid-course corrections” if and when necessary. The need to apply mid-course corrections means that contingency plans must be developed, often in advance if mid-course corrections need to be applied rapidly. The need to develop contingency plans in advance adds to the complexity of the planning process.

In practice, execution monitoring may be difficult to undertake. The attacker needs to know outcomes of various intermediate steps in the causal chain as well as what responses the victim has made at various stages of the attack, so that he may take appropriate compensating action. The difficulties of collecting such information are at least as hard as those of undertaking damage assessment for the ultimate outcome.

2.3.7 Rules of Engagement

Rules of engagement define the appropriate use of force by specifying the circumstances under which various offensive actions may be undertaken and whose authority is needed to order such actions to be taken. In the physical world, the rules of engagement may specify, for example, that individuals with guns have the authority to fire them only when they are fired upon first, and that they may never fire when the shooter is running away from them. Alternatively, rules of engagement may allow the targeting of tracked but not wheeled vehicles, or of vehicles but not personnel.

Rules of engagement specify what tools may be used to conduct a cyberattack, what targets may be attacked, what effects may be sought, and who may conduct a cyberattack under what circumstances. Rules of engagement are formulated with different intent and for different purposes depending on the interests at stake, and are discussed at greater length in Chapters 3-5 (on the military, the intelligence agencies, and law enforcement).

2.3.8 Command and Control

2.3.8.1 Command and Control—Basic Principles

According to the DOD,⁴⁸ command and control (C2) refers to the exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission. Command and control functions are performed through an arrangement of personnel, equipment, communications, facilities, and procedures employed by a commander in planning, directing, coordinating, and controlling forces and operations in the accomplishment of the mission.

In principle, C2 involves passing to weapons operators information such as the targets to be attacked, the time (which may be upon receipt) that an attack is to be launched, the nature of the attack to be launched, and the time (which may be upon receipt) that an attack in progress must be stopped. (Note, however, that not all weapons are designed to implement such C2 capabilities—sometimes, a weapon is designed to be used and then “forgotten”; that is, once launched or activated, it cannot be recalled and will do whatever damage it does without further intervention.)

C2 requires situational awareness—information about the location and status of the targets of interest and friendly or neutral entities that should not be attacked (Section 2.3.3) and their characteristics (Section 2.3.4), decision making that results in an appropriate course of action, communication of the desired course of action to the weapons available (Section 2.3.8.2, below), and damage assessment that indicates to the decision maker the results of the actions taken (Section 2.3.5).

C2 becomes more complex when more weapons, more targets, or more friendly/neutral entities must be taken into account in decision making. For example, on a complex battlefield, issues of coordination often arise. A second strike on a given target may or may not be necessary, depending on the outcome of the first strike. Target A must be attacked

⁴⁸ DOD JP1-02, *Dictionary of Military and Associated Terms*, April 12, 2001 (as amended through September 30, 2008), available at http://www.dtic.mil/doctrine/jel/new_pubs/jp1_02.pdf.

and neutralized before Target B is attacked, because Target A provides defenses for Target B. A weapon launched to attack one target may inadvertently interfere with the proper operation of another weapon launched to attack that same target—avoiding this problem is known as deconfliction. (When cyberattack is concerned, planners may well have to contend with cyberattacks launched by multiple agencies, multiple nations (e.g., allies), and even private citizens taking unauthorized actions.) All of these problems must be addressed by those planning an attack. (Section 3.6 speculates on how the DOD might address these problems.)

C2 and coordination issues are complex enough in the context of cyber activities alone. But they are multiplied enormously when cyberattacks are conducted as part of an integrated campaign (i.e., a campaign that integrates all U.S. military capabilities, not only cyber capabilities) against an adversary.

Many analysts also include matters such as damage assessment, attack assessment, and tactical warning under the rubric of command and control; this report addresses these matters in Section 2.3.5 (damage assessment), Section 2.4.1 (tactical warning and attack assessment), and Section 2.4.2 (attribution).

2.3.8.2 Illustrative Command and Control Technology for Cyberattack

A cyberattack often depends on a program running on the computer being attacked—what might be called an attack agent. The C2 function is used to convey or transmit orders about what to do, when to do it, and when to stop doing it. C2 methods can include:

- Direct (encrypted or clear text) connections to and from controller hosts or peers in peer-to-peer networks;
- Covert channels using crafted packets that appear to be innocuous protocols, controlling specific header fields in packets, or using steganographic techniques that embed commands in “normal” traffic on standard protocols (e.g., embedded characters in HTTP requests); and
- Embedded commands in files retrieved from web servers, instant messages, or chat messages on Internet Relay Chat (IRC) servers.

Direct C2 communication flows can often be detected using standard intrusion detection signature methods. Encryption may obscure the content of the information flowing in C2 channels, but it is sometimes not very hard to identify a C2 channel by looking at flow history to a host that was recently engaged in a DDOS attack or at outbound scanning activity. If a central C2 method is used, and it is easy to identify the C2 server, it can be possible to identify the attacker and to mitigate the attack.

Recently, distributed attack tool authors have sought to employ stronger cryptography, including use of public key exchange algorithms to generate per-session encryption keys, as well as to use peer-to-peer mechanisms for communication to conceal the complete distributed attack network, or to use time-delayed command execution to temporally separate C2 traffic from hostile actions like DDOS attacks that trigger alarms. Stepping stones, or relays, can further obscure the traceback from a targeted computer to the keyboards of attackers. Even as far back as 2001, the Leaves worm used both strong encryption and synchronized infected hosts' clocks to support synchronized, time-delayed execution of commands.⁴⁹

Even without centralized command and control, different attack agents can coordinate their actions. For example, an agent active on one adversary computer can delay its action until it receives a signal from a second agent on another computer that the second agent has completed its work. Or, for purposes of impeding an adversary's attempts to detect an attack, multiple agents might be implanted on a target computer with a mix of functionalities—coordination among these agents could be as effective as—or more so than—endowing a single agent with all of these functions.

Coordination among different attack agents may be particularly important if and when the same computers have been targeted by different organizations. Without careful planning, it is possible that agents may be working at cross-purposes with each other. For example, one agent may be trying to jam a communications channel that is used clandestinely by another agent.

C2 channels can also be used to update the capabilities of attack software already in place. Indeed, as long as the channel is active, such software can be upgraded or even replaced entirely—which means that an attack plan can be easily refined as more information is gained about the target. Defensive plans to prevent counterattack can also be changed in real time, so that (for example) a controller can itself be moved around (see Box 2.2).

Lastly, an attack agent will often need ways to transmit information to its controller for purposes such as damage assessment, report-back status checking, and specifying its operating environment so that a more customized attack can be put into place. For such purposes, an attacker may use outbound communications channels that are not usually blocked, such as port 80 (associated with the HTTP protocol) or DNS queries.

⁴⁹ CERT Coordination Center, "Cert Incident Note IN-2001-07 w32/leaves: Exploitation of Previously Installed Subseven Trojan Horses, July 2001. See [http://www.cert.org/incident notes/IN-2001-07.html](http://www.cert.org/incident%20notes/IN-2001-07.html).

2.3.8.3 The Role of Human Expertise and Skill

In large part because the intelligence information on a cyber target is likely to be incomplete or inaccurate in some ways, the initial stages of a cyberattack may well be unsuccessful. Thus, for a cyberattack to succeed, the attack plan may need to be modified in real time as unexpected defenses and unanticipated problems are encountered. Some cyberattacks can be “fire-and-forget”—especially those attacks for which the target set is relatively large and individual successes or failures are not particularly relevant. But if a cyberattack is very specifically targeted, adaptability and flexibility on the part of the attacker may well be needed.

2.3.9 Coordination of Cyberattack Activities with Other Institutional Entities

If a cyberattack is launched by the United States, coordination is likely to be necessary in three domains—within the U.S. government, with the private sector, and with allied nations.

- *Coordination within the U.S. government.* As noted in Chapters 3 and 4, a number of U.S. government agencies have interests in cyberattack. It is easy to imagine that a lack of interagency coordination might lead to conflicts between those wanting to exploit an adversary network and those wanting to shut it down. Policy conflicts over such matters are not new with the advent of cyberattack, but technical deconfliction issues arise as well, as different agencies might conduct cyber operations (either cyberattack and/or cyberexploitation) that might interfere with each other. In this connection, the committee has heard informally of potential struggles between the U.S. Air Force and the National Security Agency for institutional primacy in the cyberattack mission. In addition, under some circumstances, it may be necessary to consult with the congressional leadership and/or the relevant congressional committees, as discussed in Section 6.2.

- *Coordination with the private sector.* Because so much IT is designed, built, deployed, and operated by the private sector, some degree of coordination with the private sector would not be surprising in the planning and execution of certain kinds of cyberattack. For example, a cyberattack may travel over the Internet to an adversary computer, and spillover effects (such as reductions in available bandwidth) may occur that affect systems in the private sector. Or a U.S. cyberattack may prompt an adversary counterattack against U.S. systems and networks in the private sector. Or a U.S. cyberattack against an adversary transmitted through a commercial Internet service provider might be detected (and perhaps suppressed) by that provider, believing it to be the cyberattack of a crimi-

nal or acting on the protest of the targeted network. Such possibilities might suggest that the defensive posture of U.S. private sector systems and networks be strengthened in anticipation of a U.S. cyberattack (or at least that relevant commercial parties such as ISPs be notified), but this notion raises difficult questions about maintaining operational security for the U.S. attack.

- *Coordination with allied (or other) nations.* Issues of agency coordination and coordination with the private sector arise with allied nations as well, since allied nations may also have government agencies with interests in cyberattack activities and private sector entities whose defensive postures might be strengthened. Another issue is the fact that a cyberattack of the United States on Zendia might have to be transmitted over facilities controlled by third countries, and just as some nations would deny the United States military overflight rights, they may also seek to deny the United States the rights to transmit attack traffic through their facilities. Routing traffic to avoid certain countries is sometimes possible, but may require a significant amount of pre-planning and pre-positioning of assets depending on the nature of the attack to be launched.

2.3.10 A Rapidly Changing and Changeable Technology and Operational Environment for Cyberattack

The technological and operational environment in which cyberattacks may be conducted is both highly malleable and subject to very rapid change. Consider first the underlying technologies. Historical experience suggests that it takes only a decade for the technological substrate underlying IT applications to change by one, two, or three orders of magnitude—processor power, cost, bandwidth, storage, and so on. Then factor in trends of growing numbers of IT applications in both stand-alone and embedded systems and increasing connectivity among such applications. These points indicate that the overall IT environment changes on a time scale that is short compared to that of the physical world.

IT-based applications also evolve, but here the story is more mixed. Because such applications depend on knowledge and insight about how best to exploit technology, the march of progress is not nearly as consistent as it has been with the underlying technologies, and many difficult problem domains in IT have been difficult for many years. Of particular relevance to cyberattack is the problem of technical attack attribution (Section 2.4.2), which has bedeviled the cybersecurity community for many years.⁵⁰ Many cyberattack capabilities are themselves afforded by various

⁵⁰ For more discussion of this point, see National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

IT-based applications, and may—or may not—change dramatically in the future, especially in relation to the defensive capabilities available to potential victims. That is, offensive capabilities are likely to grow for all of the reasons described in Section 2.2, but defensive capabilities are also likely to grow because IT vendors are placing more emphasis on security to meet the growing criminal threat.

A second important point is that the security configuration of any given cyber target is also subject to very rapid change, and the vulnerabilities on which cyberattacks depend are sometimes easily fixed by the defender. A system administrator can close down unused access points with a few keystrokes. A patch can repair a security flaw only a few seconds after it is installed. A new security scan can discover and eliminate a malicious software agent in a few minutes. Responding to a security warning, an administrator may choose to strengthen security by deliberately degrading system functionality (e.g., reducing backward compatibility of applications that may also be associated with greater vulnerability).

Even worse from the standpoint of the attacker, all such changes in security configuration can occur without notice. (Such changes are analogous to randomly changing the schedule of a guard.) Thus, if a specific computer system is to be targeted in a cyberattack, the attacker must hope that the access paths and vulnerabilities on which the cyberattack depends are still present at the time of the attack. If they are not, the cyberattack is likely to fail.

(These considerations are less significant for a cyberattack in which the precise computers or networks attacked or compromised are not important. For example, if the intent of the cyberattack is to disable a substantial number of the desktop computer systems in a large organization, it is of little consequence that any given system is invulnerable to that attack—what matters is whether most of the systems within that organization have applied the patches, closed down unneeded access points, and so on.)

Finally, if a cyberattack weapon exploits a vulnerability that is easily closed, a change in security configuration or posture can render the weapon ineffective for subsequent use. This point is significant because it means that an attacker may be able to use a given cyberattack weapon only once or a few times before it is no longer useful. That is, certain cyberweapons may well be fragile.

2.4 CHARACTERIZING AN INCOMING CYBERATTACK

As noted in Chapter 1, the definition of active defense involves launching a cyberattack as a defensive response to an incoming cyberattack.

However, before any such response occurs, the responding party must characterize the incoming attack accurately enough that its response is both appropriate and effective. Even if the victim of an incoming cyberattack does not plan to launch a cyberattack in response, it is important to characterize the incoming attack for forensic and law enforcement purposes.

2.4.1 Tactical Warning and Attack Assessment

Tactical warning and attack assessment (TW/AA) refer to the processes through which the subject of an attack is alerted to the fact that an attack is in fact in progress and made aware of the scale, scope, and nature of an attack. In the strategic nuclear domain, early TW/AA, including, for example, information on the number of launches and their likely targets, would be provided for the United States by a network of satellites that detected adversary missiles just after launch. Moreover, the time scales of launch from Soviet territory to impact on U.S. soil (roughly 30 minutes in the case of ICBMs, 10-15 minutes in the case of submarine-launched ballistic missiles) were a primary determinant of a U.S. command and control system to assess an attack and determine the appropriate response.

For a cyberattack, even knowing that an attack is in progress is highly problematic.

- For individual sites, anomalous activity may be associated with the start of a cyberattack, and if a site detects such activity, it may receive early warning of an attack. But characterizing anomalous activity on a computer system or network that reliably indicates an attack is an enormously difficult task (legitimate users do all kinds of unusual things), and general solutions to this identification problem have eluded computer scientists for decades.

- Attack assessment is even more difficult, because the initial intrusions may simply be paving the way for hostile payloads that will be delivered later, or the damage done by a cyberattacker may not be visible for a long time after the attack has taken place (e.g., if rarely used but important data has been corrupted). (Clandestine or delayed-discovery attacks have obvious advantages when it is desirable to weaken an adversary without its knowledge.)

- A “serious” attack—that is, one conducted by a nation-state or a terrorist adversary for seriously hostile purposes—must be somehow distinguished from the background attacks that are constantly ongoing for nearly all systems connected to the Internet. These background attacks include a variety of hacking activities, virus propagation, distributed denial-of-service attacks, and other activities conducted for illicit monetary gain, sport, or pure maliciousness that are constantly being con-

ducted, in addition to the ongoing activities presumably undertaken by various nation-states or other subnational entities for covert intelligence-gathering purposes and/or to “prepare the battlefield” for possible future cyberattacks for offensive purposes.

For a dispersed entity (such as the Department of Defense, the U.S. government, or a large corporation), multiple sites may be attacked in a coordinated manner. If attacks were somehow known to be coordinated, such coordination might indicate a serious attack. On the other hand, detecting such coordination against the background noise of ongoing attacks also remains an enormous intellectual challenge, as useful information from multiple sites must be made available on a timely basis. (And as detection capabilities improve, attackers will take steps to mask such signs of coordinated attacks.)

An attack assessment would seek to address many factors, including the scale of the attack (how many entities are under attack), the nature of the targets (which entities are under attack, e.g., the DOD Global Command and Control System, electric power generating facilities, Internet retailers), the success of the attack and the extent and nature of damage caused by the attack, the extent and nature of any foreign involvement derived from technical analysis of the attack and/or any available intelligence information not specifically derived from the attack itself, and attribution of the source of the attack (discussed at greater length in Section 2.4.2).

Information on these factors is likely to be quite scarce when the initial stages of an attack are first noticed. For example, because cyberweapons can act over many time scales, anonymously, and clandestinely, knowledge about the scope and character of a cyberattack will be hard to obtain quickly. Other non-technical factors may well play into an attack assessment, such as the state of political relations with other nations that are capable of launching such an attack.

From an organizational perspective, the response of the United States to a cyberattack by a non-state actor is often characterized as depending strongly on whether the attack—as characterized by factors such as those described above—is one that requires a law enforcement response *or* a national security response. This characterization is based on the idea that a national security response relaxes many of the constraints that would otherwise be imposed by a law enforcement response.⁵¹

But the “law enforcement versus national security” dichotomy is

⁵¹ For example, active defense—either by active threat neutralization or by cyber retaliation—may be more viable under a national security response paradigm, whereas a law enforcement paradigm might call for passive defense to mitigate the immediate threat and other activities to identify and prosecute the perpetrators.

misleading. In practice, the first indications of a cyberattack are likely to be uncertain, and many factors relevant to a decision will be unknown. Once the possibility of a cyberattack is made known to national authorities, information must be gathered to determine perpetrator and purpose, and must be gathered using the available legal authorities (described in Section 7.3). Some entity within the federal government integrates the relevant information and then it or another higher entity (e.g., the National Security Council) renders a decision about next steps to be taken, and in particular whether a law enforcement or national security response is called for.

How might some of the factors described above be taken into account as a greater understanding of the event occurs? Law enforcement equities are likely to predominate in the decision-making calculus if the scale of the attack is small, if the assets targeted are not important military assets or elements of critical infrastructure, or if the attack has not created substantial damage. To the extent that any of these characteristics are not true, pressures may increase to regard the event as one that also includes national security equities.

The entity responsible for integrating the available information and recommending next steps to be taken has evolved over time. In the late 1990s, the U.S. government established the National Infrastructure Protection Center (NIPC) as a joint government and private sector partnership that provided assessment, warning, vulnerability, and investigation and response for threats to national critical infrastructure. Consisting of personnel from the law enforcement, defense, and intelligence communities, each with reach-back into their respective agencies for support, along with representatives from the private sector and foreign security agencies, the NIPC was the place where information on the factors described was to be fused and the intelligence, national security, law enforcement, and private sector equities integrated regarding the significance of any given cyberattack.

Organizationally, the NIPC was part of the Department of Justice under the Federal Bureau of Investigation. In later years, the analysis and warning functions of the NIPC were dispersed throughout the Department of Homeland Security (DHS) as the result of that department's creation, while the principal investigative functions remained at the FBI (with some investigative functions performed by the U.S. Secret Service, an autonomous part of DHS).⁵² Initially, they were integrated into the Information Analysis and Infrastructure Protection Directorate, primarily the National Infrastructure Coordinating Center (NICC) under the Office

⁵² See Department of Homeland Security, "History: Who Became Part of the Department?," 2007, available at http://www.dhs.gov/xabout/history/editorial_0133.shtm.

of Operations Coordination and the United States Computer Emergency Readiness Team, the operational arm of the National Cyber Security Division. The NICC provides operational assessment, monitoring, coordination and response activities, and information sharing with the private sector through information sharing and analysis centers. The United States Computer Emergency Readiness Team (US-CERT), the operational arm of the National Cyber Security Division, coordinates defense against cyberattacks.

Further reorganization at DHS moved the Office of Operations Coordination to a freestanding component that runs NICC as part of the National Operations Center. The Office of Infrastructure Protection (OIP) became part of the National Protection and Programs (NPP) Directorate. A separate Office of Cybersecurity and Communications, also under NPP, includes the National Cyber Security Division, which still manages US-CERT operations. Broadly, the NIPC functions that focused on risk reduction, warning, and vulnerability assessment are now part of NPP. Those NIPC functions that focused on operational assessment and coordination are today part of the NICC under the Office of Operations Coordination.

As this report is being written, the U.S. government apparatus responsible for warning and attack assessment is likely to be reorganized again.

The government agencies responsible for threat warning and attack assessment can, in principle, draw on a wide range of information sources, both inside and outside the government. In addition to hearing from private sector entities that come under attack, cognizant government agencies can communicate with security IT vendors, such as Symantec and McAfee, that monitor the Internet for signs of cyberattack activity. Other public interest groups, such as the Open Net Initiative and the Information Warfare Monitor, seek to monitor cyberattacks launched on the Internet.⁵³

2.4.2 Attribution

Attribution is the effort to identify the party responsible for a cyber-attack. *Technical attribution* is the ability to associate an attack with a responsible party through technical means based on information made available by the fact of the cyberattack itself—that is, technical attribution

⁵³ See <http://opennet.net/> and <http://www.infowar-monitor.net> for more information on these groups. A useful press report on the activities of these groups can be found at Kim Hart, "A New Breed of Hackers Tracks Online Acts of War," *Washington Post*, August 27, 2008, available at http://www.washingtonpost.com/wp-dyn/content/article/2008/08/26/AR2008082603128_pf.html.

is based on the clues available at the scene (or scenes) of the attack. All-source attribution is a process that integrates information from all sources, not just technical sources at the scene of the attack, in order to arrive at a judgment (rather than a definitive and certain proof) concerning the identity of the attacker.

Two key issues in technical attribution are precision and accuracy:

- *Precision.* An attribution has associated with it some range of precision for the identity of the attacker. The attack might be associated with a specific nation (Zendia), a specific department within that nation (the ministry of defense), a specific unit (the 409th Information Operations Brigade), a specific set of IP addresses, a specific individual, and so on.

- *Accuracy.* A characteristic related to precision is accuracy, a measure of the quality of attribution, such as the probability that the attribution is correct. Accuracy is a key issue in legal standards for evidence and in the extent to which it is reasonable to develop linkages and inferences based on those attributes. Note that an attacker may seek to reduce the accuracy of attribution if he or she wishes to operate secretly by taking countermeasures to impersonate other parties.

The unfortunate reality is that technical attribution of a cyberattack is very difficult to do (it is often said that “electrons don’t wear uniforms”), and can be nearly impossible to do when an unwittingly compromised or duped user is involved. As the existence of botnets illustrates, a cyberattacker has many incentives to compromise others into doing his or her dirty work, and untangling the nature of such a compromise is inevitably a time-consuming and laborious (if not futile) process.

To illustrate the point, consider a scenario in which computers of the U.S. government are under a computer network attack (e.g., as the result of a botnet attack). The owners/operators of the attacked computers in the United States would likely be able to find the proximate source(s) of any attack. They might discover, for example, that the attack traffic emanated from computers located in Zendia. But there may well be no technical way to differentiate among a number of different scenarios consistent with this discovery. These scenarios include the following:

- The attack against the United States was launched by agents of the Zendian government with the approval of the Zendian national command authority.
- The attack against the United States was launched by low-level agents of the Zendian government without the approval or even the knowledge of the Zendian national command authority.
- The attack was launched through the efforts of computer-savvy

citizens of Zendia who believe that the United States oppresses Zendia in some way. Although the efforts of these citizens are not initiated by the Zendian government, the Zendian government takes no action to stop them. (Such individuals are often known as “patriotic hackers” and are discussed in more detail in Section 7.2.3.3.)

- The Zendian computers used to conduct the attack against the United States have been compromised by parties outside Zendia (perhaps even from the United States, as happened in the Solar Sunrise incident in February 1998⁵⁴), and Zendia is merely an innocent bystander on the international stage.

- The attack was launched at the behest of the Zendian government, but not carried out by agents of the Zendian government. For example, it may have been carried out by the Zendian section of an international criminal organization.

However, the limitations of technical attribution are not dispositive. All-source attribution takes into account whatever information is available from efforts at technical attribution, but also uses information from other sources to arrive at a judgment. Such sources might include:

- *Intelligence sources.* For example, a well-placed informant in the Zendian government might provide information indicating the responsibility of that nation in initiating the attack, or routinely monitored message traffic might indicate a point of responsibility within the Zendian government.

- *Political sources.* The Zendian government might publicly take credit for the attack. (Of course, a claim that “We were responsible for the attack” would itself have to be verified.)

- *Other technical information.* The technical signature of the cyberattack might be similar to that of a previous attack, and the United States might have information on the originator of that previous attack. The scale or nature of the attack might be such that only a major nation-state could have mounted it, thus ruling out other parties. Or it might be possible to determine the time zone of the actual attacking machine.⁵⁵

- *Temporal proximity to other coercive or aggressive actions that can be attributed.* For example, Zendia might choose to “bundle” a set of such actions together, such as cyberattack coupled with an embargo on selling

⁵⁴ More information on the Solar Sunrise incident can be found at http://www.sans.org/resources/idfaq/solar_sunrise.php.

⁵⁵ For example, it is sometimes possible to learn information about a target computer’s physical environment through the remote monitoring of time stamps. Local time stamps are governed by a computer’s clock, and the rate at which the clock runs is affected by the ambient temperature. Thus, time stamp information provides information on changes of ambient

certain computer chips or strategic raw materials to the United States, a break in diplomatic relations, and refusal of “safe harbor” rights for U.S. naval vessels.

Thus, although the process of all-source attribution might well take a long time to arrive at an actionable (if not definitive) judgment, the case for attribution is not as hopeless as it is often portrayed.

Attribution of an attack should not be confused with establishing or identifying an access path to the source of the attack. Under a given set of circumstances, the victim may be able to establish both of these pieces of information, one of them, or none of them. For example, it may be impossible to establish an access path to the source of a cyberattack, but at the same time an all-source attribution effort might definitively identify a given nation as being responsible for the attack. Alternatively, an access path to the source of a cyberattack might be established without providing any useful information at all regarding the party responsible (e.g., the launching point for a cyberattack against a corporation might be located inside that corporation and not be further traceable). The difference between attribution and having an access path is significant, because in the absence of an access path, neutralization of a cyberattack is not possible, though retaliation for it might be. The converse is true as well—in the absence of attribution, retaliation or reprisal is not possible, though neutralization of a cyberattack might be.

Finally, the problems that anonymity poses for a defensive actor can easily be seen as advantages for an attacker. The discussion above suggests that with careful and appropriate planning, a cyberattack can often be conducted with a high degree of plausible deniability. Such a capability may be useful in certain intelligence operations when it is desirable that the role of the government sponsor of a cyberattack is not to be publicly acknowledged (as discussed in Section 4.2).

2.4.3 Intent

In the realm of traditional military conflict, it is generally presumed that national governments control the weapons of warfare—frigates,

temperature, which may be correlated with time-of-day physical location. Measurements of day length and time zone can provide information useful for estimating the physical location of a computer. Local temperature changes caused by air-conditioning or movements of people can identify whether two machines are in the same location, or even are virtual machines on one server. See Steven J. Murdoch, “Hot or Not: Revealing Hidden Services by Their Clock Skew,” Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS’06, October 30–November 3, 2006, Alexandria, Va., 2006, available at <http://www.cl.cam.ac.uk/users/sjm217/>.

fighter jets, tanks, and so on. Thus, if any of these weapons are used, there is a presumption that actions involving them have been sanctioned by the controlling government—and inferences can often be drawn regarding that government's intent in ordering those actions.

But when other weapons are not controlled exclusively by governments, inferring intent from action is much more problematic. This is especially so if communication cannot be established with the controlling party—as will often be the case with cyberattack. Attribution of a cyberattack (discussed above) helps, but if the party identified as being responsible is not a national government or another party with declared intentions toward the United States, it will be virtually impossible to determine intent with high confidence.

Determinations of intent and attribution of the source are often complicated—and inappropriately biased—by a lack of information. Ultimately, such determinations are made by human beings, who seek to integrate all available information in order to render a judgment. (Such integration may be automated, but human beings program the rules for integration.) When inexperienced human beings with little hard information are placed into unfamiliar situations in a general environment of tension, they will often make worst-case assessments. In the words of a former Justice Department official involved with critical infrastructure protection, “I have seen too many situations where government officials claimed a high degree of confidence as to the source, intent, and scope of an attack, and it turned out they were wrong on every aspect of it. That is, they were often wrong, but never in doubt.”

2.5 ACTIVE DEFENSE FOR NEUTRALIZATION AS A PARTIALLY WORKED EXAMPLE

To suggest how the elements above might fit together operationally, consider how a specific active defense scenario might unfold. In this scenario, active defense means offensive actions (a cyber counterattack) taken to neutralize an immediate cyberthreat—that is, with an operational goal—rather than retaliation with a strategic goal. The hostile cyberattack serves the offensive purposes of Zendia. The cyber counterattack in question is for defensive purposes.

The scenario begins with a number of important U.S. computer systems coming under cyberattack. For definiteness, assume that these computer systems are SCADA and energy management systems controlling elements of the power grid, and that the attacker is using unauthorized connections between these systems and the Internet-connected business systems of a

power generation facility to explore and manipulate the SCADA and energy management systems.

The first step—very difficult in practice—is to recognize the act as an unambiguously hostile one rather than one undertaken by cyber pranksters with too much time on their hands. Further inspection reveals that the unauthorized intruder has planted software agents that would respond to commands in the future by disabling some of the power generation and transmission hardware controlled by these systems, and furthermore that the apparent controllers of these agents are located around the world. However, even the availability of such information cannot determine the motivations of the responsible parties regarding why they are undertaking such a hostile act.

A second step is to recognize that these attacks are occurring on many different SCADA and energy management systems around the nation. Such recognition depends on the existence of mechanisms within the U.S. government for fusing information from different sources into an overall picture indicating that any individual attack fits into a larger adversarial picture, rather than being an isolated event.

The third step is to identify the attacker—that is, the party installing the agents. The IP address of the proximate source of this party can be ascertained with some degree of confidence, and a corresponding geographic location may be available—in this case, the geographic location of the proximate source is Zendia. But these facts do not reveal whether the attack was:

- Sponsored by Zendia and launched with the approval of the highest levels of the Zendian National Command Authority;
- Launched by low-level elements in the Zendian military without high-level authorization or even the knowledge of the Zendian NCA;
- Launched by computer-savvy Zendian citizens;
- Launched by terrorists from Zendian soil; or
- Launched by Ruritania transiting through Zendia, which may be entirely innocent.

Suppose further that additional information from non-technical sources is available that sheds additional light on the attacker's identity. In this case, intelligence sources indicate with a moderate degree of confidence that the attack ultimately emanates from parties in Zendia.

The availability of information about the attacker's identity marks an important decision point about what to do next. One option is to approach the Zendian government and attempt to resolve the problem diplomatically and legally, where "resolution" would call for Zendian government action that results in a cessation of the attack—in this case, refraining from install-

ing any more agents on U.S. SCADA and energy management systems. (Knowing that the attacks have actually ceased is yet another problem, especially against the background of myriad other hostile or adversarial actions being taken every day against U.S. systems of various sorts.) Such an approach also risks compromising U.S. intelligence sources, and thus U.S. decision makers may be wary of taking this route.

Continuing with this scenario, the United States discovers that the hostile agent controllers are themselves centrally controlled by an Internet-connected system located in Zendia. Cognizant of the uncertainties involved, the United States quietly probes the master controller to understand its vulnerabilities, but decides to refrain from further action at this time. It also works on removing the deployed agents from the SCADA and energy management systems in question, replacing them with harmless look-alike agents that can perform all of the relevant report-back functions to the controller. However, cyber response teams from the United States realize that they are unlikely to find every SCADA and energy management system so infested.

A few months later, tensions between the United States and Zendia rise because of a non-lethal incident between the Zendian air force and a U.S. reconnaissance plane. In order to put pressure on the United States, Zendia tries to activate its SCADA/EMS agents. Zendia receives many affirmative reports from the field, some of which are in fact misleading and others valid. In order to stop the remaining agents, the United States launches a denial-of-service attack against the Zendian controller, effectively disconnecting it from the Internet while at the same time issuing a demarche to the Zendian government to cease its hostile actions and to provide information on the SCADA/EMS systems penetrated that is sufficient to effect the removal of all hostile agents. Zendia responds to the U.S. demarche publicly, denouncing the U.S. denial-of-service attack on it as an unprovoked and unwarranted escalation of the situation.

This neutralization scenario raises many issues. Neutralization of cyberthreats requires an access path to the particular hardware unit from which the attack emanates (e.g., the attack controller). In addition, an indication of the physical location of that hardware may be necessary.

- Knowledge of the controller's specific hardware unit is important because the attacker may have taken a very circuitous route to reach the target. If the attacker has been clever, neutralization of any intermediate node along the way is unlikely to result in a long-term cessation of the attack, and only disruption of the controller will suffice. If not, there may be a particular intermediate node whose destruction or degradation may be sufficient to stop the attack.

- Physical location is important because of the legal jurisdictional issue—depending on the physical (national) location of the hardware, different laws regarding the putative criminality of its behavior and the legality of damaging it may apply. This point is relevant especially if an attack has a foreign origin; however, knowledge of physical location is not *required* to neutralize the attack.

In practice, none of these conditions are easy to meet. Attackers have strong incentives to conceal their identity, and so are likely to use compromised computers as intermediate launching points for their attacks. Furthermore, because one compromised computer can be used to compromise another one, the chain leading to the actual attacker—the only one with malevolent intent—can be quite long, and thus quite difficult (and time-consuming) to unravel. By the time an actual machine identity of the controller has been established, the attacker may no longer have a presence on the originating machine.

Yet another complicating factor is that the controller function can be executed on a variety of different systems. Thus, even if a victim is successful in identifying the controller of the attack, and even if it successfully launches a counterattack that neutralizes that controller (a counterattack that may be electronic, kinetic, or even legal), the controller function may shift automatically to another system—if so, another laborious process may need to be started again. (An analogy could be drawn to the operation of a mobile missile launcher [Transporter-Erector-Launcher, TEL]. A TEL sets down in a specific, usually pre-surveyed, location, launches its missile, and then immediately moves to minimize the effectiveness of a counterattack against it.) On the other hand, the controller of the attack may not shift, especially if the attacker is not well resourced or sophisticated. Under such circumstances, a counter-cyberattack may well succeed in shutting down an attack, at least for a while.

A long chain of compromised machines is not the only obfuscation technique an attacker may use. An attacker may also plant false evidence implicating other parties on one or more of the intermediate links. Such evidence could lead the forensic investigator to mistakenly identify a particular intermediate node as the true source of an attack, and a neutralization counterattack launched against that node would target an innocent party. In this case, the fact that the United States has only moderate confidence in the fact of Zendian responsibility is problematic.

An important aspect of any neutralization counterattack is the time it takes to determine the identity of the attacking party and to establish an access path and its geographic location. Perhaps the most plausible justification for a neutralization counterattack is that a counterattack is

needed to stop the immediate harm being done by an attack. If it takes so long for the defending party to obtain the necessary information for a counterattack that the attack has ceased by the time the counterattack can be launched, this justification may no longer be plausible.⁵⁶

Note that the policy requirement to quickly and properly identify the attacking party and the technical reality that attribution is a time-consuming task work against each other. Shortening the time for investigation, perhaps by going so far as to automate the identification, assessment, and response process, may well increase the likelihood of errors being made in any response (e.g., responding against the wrong machine, launching a response that has large unintended effects).

On the other hand, it is possible that a neutralization cyberattack would be used only after a number of hostile cyberattacks had occurred. Consider the ease of an unknown adversary launching cyberattacks against a particular U.S. defense facility. If forensic investigation is undertaken after each attack, after a while enough information might be obtained to determine the leading indicators of an attack by this adversary. At some point, the United States might well have enough information in hand so that it could respond quickly to the next cyberattack launched by this adversary, and might be willing to take the chance that it was responding erroneously with a neutralization cyberattack.

From a policy standpoint, the acceptability of an increase in the likelihood of errors almost surely depends on the state of the world at the time. During times of normal political relations with other nations, such an increase may be entirely unacceptable. However, during times of political, diplomatic, or even military tension with other nations, the U.S. leadership might well be willing to run the risk of a mistaken response in order to ensure that a response was not crippled by an adversary attack. (In this regard, the situation is almost exactly parallel to the issue of riding out a strategic attack on the United States or employing a strategy of launching a land-based strategic missile on warning or while under attack—the latter being regarded as much more likely during times of tension with a putative adversary.)

Under some circumstances, the United States might choose to launch a neutralization cyberattack fully expecting that the adversary would respond with an even larger hostile cyberattack. If it did so, it would be necessary for the United States to prepare for that eventuality. Such preparation might involve taking special measures to strengthen the cybersecurity posture of key facilities and/or preparing for kinetic escalation.

⁵⁶ On the other hand, the cessation of an attack may simply indicate the end of one phase and the start of a lull before the next phase. A clever attacker would launch the next phase in such a way that the defender would have to unravel an entirely new chain.

These concerns do not automatically imply that neutralization counterattacks are a bad idea under all circumstances. But they do raise several questions that must be answered before such a response is made.

- What defensive measures must be taken, if any, before launching a neutralization counterattack? Should a neutralization counterattack be a last resort, to be used when all other methods for responding to a cyberattack have proven (or will prove) ineffective?⁵⁷ Or should a neutralization counterattack be a measure of first resort, to be triggered automatically without human intervention in the first few seconds of an attack? Or somewhere in between?
- A counterattack requires only that an access path to the attacker be available. Under what circumstances must the identity of the attacking party be known? If the attacker must be known, what degree of confidence and what evidentiary basis are needed? And how, if at all, should the attacker's identity affect a decision to launch a counterattack? (For example, how might such a decision be affected by the fact that an attack is emanating from the information network of a U.S. hospital or an important laboratory?)
- How likely is it that the attacker will have anticipated a neutralization counterattack and taken steps to mitigate or negate the effect of the counterattack? What are those steps likely to have been? How likely is it that a neutralization counterattack will indeed curb or halt the incoming attack?
- How narrowly should a neutralization counterattack be focused? Should it be limited solely to eliminating or mitigating the threat (and not causing harm outside that effort)? Or is causing additional damage to the attacker a desirable outcome?
- At what threshold of actual or expected damage to U.S. systems and networks should a neutralization counterattack be launched? That is, how should the benefit of a counterattack be weighed against the political risks of launching it? For example, what targets are worth protecting? (U.S. military installations? Installations associated with national critical infrastructure? Defense industrial base firms? Fortune 500 companies?)
- How should the threshold of damage be established? Should it be established unilaterally in real time by the original victim (e.g., the corporation or government entity attacked)? Or should it result from an orderly interagency and governmental process that operates well in advance of when policy guidance is needed?

⁵⁷ For example, one might argue that technical means such as target hardening and adversary deception and legal methods such as appeal to an ISP to disconnect an attacker from the Internet must be exhausted before active defense is considered.

BOX 2.4 A Possible Taxonomy of Active Responses

There is a broad range of actions possible to respond to a cyberattack. One possible taxonomy of response actions was developed by Sergio Caltagirone.¹ This taxonomy identifies eight types of response in increasing order of activity required by the responder, potential impact on the attacker, and potential for collateral and unintended consequences.

1. *No action*—a conscious decision to take no action in response to an identified attack. Not taking any action is active insofar as it involves a thoughtful decision process that considers the benefits and costs of potential options.

2. *Internal notification*—notifying users, administrators, and management of the system attacked. Some subset of these may be notified depending on the type of attack, but the attack is not reported to anyone outside the organization of the affected system.

3. *Internal response*—taking specific action to protect the system from the attacker. The response likely depends on the type of attack, but might include blocking a range of IP addresses or specific ports, segmenting or disconnecting parts of the system, and purposely dropping connections.

4. *External cooperative response*—contacting external groups or agencies with responsibility for classifying, publicizing and analyzing attacks (e.g., CERT, DShield), taking law enforcement action (e.g., FBI, Secret Service), providing protection services (e.g., Symantec, MacAfee), and providing upstream support (e.g., Internet service providers).

There is a broad consensus that Actions 1-4 are legitimate actions under almost any set of circumstances. That is, an individual or organization is unambiguously allowed to take any of these actions in response to a cyberattack. However, the same is not true for Actions 5-8 described below, which are listed in order of increasing controversy and increasing likelihood of running afoul of today's legal regime should the target of a cyberattack take any of these actions.

Lastly, given the difficulties of knowing if a cyberattack is taking or has taken place; whether a given cyberattack is hostile, criminal, or mischievous in intent; the identity of the responsible party; and the extent to which it poses a significant threat, the neutralization option must not be seen as the only way to respond to an attack. Box 2.4 describes a spectrum of possible responses to a cyberattack—note that the neutralization option corresponds to Action 6 or Action 7, and as such is a more aggressive form of response.

5. *Non-cooperative intelligence gathering*—the use of any tools to gather information about the attack and the attacker. Tools might include honeypots, honeynets, traceroutes, loose source and record routes, pings and fingers.

6. *Non-cooperative “cease and desist”*—the use of tools to disable harmful services on the attacker’s system without affecting other system services.

7. *Counterstrike*—response taking two potential forms: direct action (active counterstrike) such as hacking the attacker’s systems (hack-back) and transmitting a worm targeted at the attacker’s system; passive counterstrike that redirects the attack back to the attacker, rather than directly opposing the attack. Examples of passive counterstrike are a footprinting strike-back that sends endless data, bad data, or bad SQL requests, and network reconnaissance strike-back using trace-route packets (ICMP “TTL expired”).

8. *Preemptive defense*—conducting an attack on a system or network in anticipation of that system or network conducting an attack on your system.

Different actions may be taken based on the type of attack and an analysis of the benefits and costs associated with each type of response. Multiple types of responses may be taken for any given attack.

Actions 1-4 are generally non-controversial, in the sense that it would not be legally problematic for a private company to take any of these responses. Actions 6-8 are much more aggressive, fall into the general category of active defense (and more), and certainly raise many questions under the statutory prohibitions against conducting cyberattack. In addition, system administrators often express concern about the legality of Action 5 in light of the various statutes governing electronic surveillance.

¹ S. Caltagirone and D. Frincke, *Information Assurance Workshop, 2005, IAW '05, Proceedings from the Sixth Annual IEEE SMC*, June 15-17, 2005, pp. 258-265. See also David Dittrich and Kenneth Einar Himma, “Active Response to Computer Intrusions,” *The Handbook of Information Security*, Hossein Bidgoli, editor-in-chief, John Wiley & Sons, Inc., Hoboken, N.J., 2005.

2.6 TECHNICAL AND OPERATIONAL CONSIDERATIONS FOR CYBEREXPLOITATION

2.6.1 Technical Similarities in and Differences Between Cyberattack and Cyberexploitation

The cyberexploitation mission is different from the cyberattack mission in its objectives (as noted in Chapter 1) and in the legal constructs surrounding it (as discussed in Chapter 7). Nevertheless, much of the technology underlying cyberexploitation is similar to that of cyberattack, and the same is true for some of the operational considerations as well.

As noted in Section 2.2.2, a successful cyberattack requires a vulnerability, access to that vulnerability, and a payload to be executed. A cyber-exploitation requires the same three things—and the only technological difference is in the payload to be executed. That is, what distinguishes a cyberexploitation from a cyberattack is the nature of the payload.

Whereas the attacker might destroy the papers inside a locked file cabinet once he gains access to it, the exploiter might copy them and take them away with him. In the cyber context, the cyberexploiter will seek to compromise the confidentiality of protected information afforded by a computer system or network.

2.6.2 Possible Objectives of Cyberexploitation

What might cyberexploitations seek to accomplish? Here are some hypothetical examples. The cyberexploiter might seek to:

- *Exploit information available on a network.* For example, an attacker might monitor passing traffic for keywords such as “nuclear” or “plutonium,” and copy and forward to the attacker’s intelligence services any messages containing such words for further analysis. A cyberexploitation against a military network might seek to exfiltrate confidential data indicating orders of battle, operational plans, and so on. Alternatively, passwords are often sent in the clear through e-mail, and those passwords can be used to penetrate other systems. This objective is essentially the same as that for all signals intelligence activities—to obtain intelligence information on an adversary’s intentions and capabilities.

- *Be a passive observer of a network’s topology and traffic.* As long as the attacker is a passive observer, the targeted adversary will experience little or no direct degradation in service or functionality offered by the network. Networks can be passively monitored to identify active hosts as well as to determine the operating system and/or service versions (through signatures in protocol headers, the way sequence numbers are generated, and so on).⁵⁸ The attacker can map the network and make inferences about important and less important nodes on it simply by performing traffic analysis. (What is the organizational structure? Who holds positions of authority?) Such information can be used subsequently to disrupt the network’s operational functionality. If the attacker is able to read the contents of traffic (which is likely, if the adversary believes the network is secure and thus has not gone to the trouble of encrypting

⁵⁸ Annie De Montigny-Leboeuf and Frederic Massicotte, “Passive Network Discovery for Real Time Situation Awareness,” 2004, available at <http://www.snort.org/docs/industry/ADeMontignyNatoISToulouse2004.pdf>.

traffic), he can gain much more information about matters of significance to the network's operators. As importantly, a map of the network provides useful information for a cyberattacker, who can use this information to perform a more precise targeting of later attacks on hosts on the local network, which are typically behind firewalls and intrusion detection/prevention systems that might trigger alarms.

- Obtain technical information from a company's network in another country in order to benefit a domestic competitor of that company. For example, two former directors of the DGSE (the French intelligence service) have publicly stated that one of the DGSE's top priorities was to collect economic intelligence. During a September 1991 NBC news program, Pierre Marion, former DGSE director, revealed that he had initiated an espionage program against U.S. businesses for the purpose of keeping France internationally competitive. Marion justified these actions on the grounds that the United States and France, although political and military allies, are economic and technological competitors. During an interview in March 1993, then-DGSE director Charles Silberzahn stated that political espionage was no longer a real priority for France but that France was interested in economic intelligence, "a field which is crucial to the world's evolution." Silberzahn advised that the French have had some success in economic intelligence but stated that much work is still needed because of the growing global economy. Silberzahn advised during a subsequent interview that theft of classified information, as well as information about large corporations, was a long-term French government policy.⁵⁹

The examples above suggest certain technical desiderata for cyberexploitations. For instance, it is highly desirable for a cyberexploitation to have a signature that is difficult for its target to detect, since the cyberexploitation operation may involve many separate actions spread out over a long period of time in which only small things happen with each action. One reason is that if the targeted party does not know that its secret information has been revealed, it is less likely to take countermeasures to negate the compromise. A second reason is that the exploiter would like to use one penetration of an adversary's computer or network to result in multiple exfiltrations of intelligence information over the course of the entire operation. That is, the intelligence collectors need to be able to maintain a clandestine presence on the adversary computer or network despite the fact that information exfiltrations provide the adversary with opportunities to discover that presence.

Also, an individual payload can have multiple functions simultane-

⁵⁹ See page 33, footnote 1, in National Research Council, *Cryptography's Role in Securing the Information Society*, National Academy Press, Washington, D.C., 1996.

ously—one for cyberattack and one for cyberexploitation—and which function is activated at any given time will depend on the necessary command and control arrangements (see Section 2.3.8). For example, a payload delivered to an adversary command and control network may be designed to exfiltrate information during the initial stages of a conflict and then to degrade service on the network when it receives a command to do so.

In addition, the relationship between technologies for cyberexploitation and cyberattack is strong enough that the cost of equipping a tool for the former with the capability for the latter is likely to be low—so low that in many cases acquisition managers could find it sensible as a matter of routine practice to equip a cyberexploitation tool with attack capabilities (or provide it with the ability to be modified on-the-fly in actual use to have such capabilities).⁶⁰

2.6.3 Approaches for Cyberexploitation

As is true for cyberattack, cyberexploitation can be accomplished through both remote-access and close-access methodologies.

A hypothetical example of cyberexploitation based on remote access might involve “pharming” against an unprotected DNS server, such as the one resident in wireless routers.⁶¹ Because wireless routers at home tend to be less well protected than institutional routers, they are easier to compromise. Successful pharming would mean that web traffic originating at the home of the targeted individual (who might be a senior official in an adversary’s political leadership) could be redirected to websites controlled by the exploiter. With access to the target’s home computer thus provided, vulnerabilities in that computer could be used to insert a payload that would exfiltrate the contents of the individual’s hard disk, possibly providing the exploiter with information useful for blackmailing the target. As a historical precedent, Symantec in January 2008 reported an incident directed against a Mexican bank in which the DNS settings on a customer’s home router were compromised.⁶² An e-mail was sent to the target, ostensibly from a legitimate card company. However, the e-mail

⁶⁰ If these cyberexploitation tools were to be used against U.S. citizens (more precisely, U.S. persons as defined in EO 12333 (Section 7.3.6)), legal and/or policy implications might arise if these tools were to have attack capabilities as well. Thus, the observation is most likely to be true for tools that are not intended for such use.

⁶¹ “Pharming” is the term given to an attack that seeks to redirect the traffic to a particular website to another, bogus website.

⁶² Ellen Messmer, “First Case of ‘Drive-by Pharming’ Identified in the Wild,” *Network World*, January 22, 2008, available at <http://www.networkworld.com/news/2008/012208-drive-by-pharming.html>.

contained a request to the home router to tamper with its DNS settings. Thus, traffic intended for the bank was redirected to the criminal's web-site mimicking the bank site.

A hypothetical example of cyberexploitation based on close access might involve intercepting desktop computers in their original shipping cartons while they are awaiting delivery to the victim, and substituting for the original video card a modified one that performs all of the original functions and also monitors the data being displayed for subsequent transmission to the exploiter. There is historical precedent for such approaches. One episode is the 1984 U.S. discovery of Soviet listening devices in the Moscow embassy's typewriters—these devices captured all keystrokes and transmitted them to a nearby listening post.⁶³ A second reported episode involves cameras installed inside Xerox copiers in Soviet embassies in the 1960s.⁶⁴ A third episode, still not fully understood, is the 2004-2005 phone-tapping affair in Greece.⁶⁵

2.6.4 Some Operational Considerations for Cyberexploitation

2.6.4.1 The Fundamental Similarity Between Cyberattack and Cyberexploitation

Because the cyber offensive actions needed to carry out a cyber-exploitation are so similar to those needed for cyberattack, cyberexploitations and cyberattacks may be difficult to distinguish in an operational context. (The problem of distinguishing between them is compounded by the fact that an agent for exploitation can also contain functionality to be used at another time for attack purposes.) This fundamental ambiguity—absent with kinetic, nuclear, biological, and chemical weapons—has several consequences:

⁶³ Jay Peterzell, "The Moscow Bug Hunt," *Time*, July 10, 1989, available at <http://www.time.com/time/magazine/article/0,9171,958127-4,00.html>.

⁶⁴ Ron Laytner, "Xerox Helped Win The Cold War," *Edit International*, 2006, available at <http://www.editinternational.com/read.php?id=47ddf19823b89>.

⁶⁵ In this incident, a number of mobile phones belonging mostly to members of the Greek government and top-ranking civil servants were found to have been tapped for an extended period of time. These individuals were subscribers to Vodafone Greece, the country's largest cellular service provider. The taps were implemented through a feature built into the company's switching infrastructure originally designed to allow law enforcement agencies to tap telephone calls carried on that infrastructure. However, those responsible for the taps assumed control of this feature to serve their own purposes and were also able to conceal their activities for a long time. The sophistication of the programming required to undertake this compromise is considerable, and has led to speculation that the affair was the result of an inside job. See Vassilis Prevelakis and Diomidis Spinellis, "The Athens Affair," *IEEE Spectrum*, July 2007, available at <http://www.spectrum.ieee.org/print/5280>.

- The targeted party may not be able to distinguish between a cyber-exploitation and a cyberattack, especially on short time scales, even if such differences are prominent in the minds of the party undertaking cyber offensive actions.
- Because the legal authorities to conduct cyberexploitations and cyberattacks are quite different, clarity in the minds of the operators about their roles in any given instance is essential.
- From a training and personnel standpoint, developing expertise at cyberattack also develops most of the required skill set for conducting cyberexploitation, and vice versa.⁶⁶

2.6.4.2 Target Identification and Intelligence Preparation

Although some intelligence operations may be characterized by a “vacuum cleaner” approach that seeks to obtain all available traffic for later analysis, a cyberexploiter may be very concerned about which computers or networks are targeted—an issue of precision. Very precise cyber-exploitations would be characterized by small-scale operations against a specific computer or user whose individual compromise would have enormous value (“going after the crown jewels”)—the vice president’s laptop, for example.

To the extent that specific systems must be targeted, substantial intelligence efforts may be required to identify both access paths and vulnerabilities. For example, even if the vice president’s laptop is known to be a Macintosh running OS-X, there may well be special security software running on her laptop; finding out even what software might be running, to say nothing of how to circumvent it, is likely to be very difficult in the absence of close access to it. The same considerations are true of Internet-connected computer systems that provide critical functionality to important companies and organizations—they may well be better pro-

⁶⁶ For example, Air Force Doctrine Document 2-5 (issued by the Secretary of the Air Force, January 11, 2005) explicitly notes that “military forces under a combatant commander derive authority to conduct NetA [network attack] from the laws contained in Title 10 of the U.S. Code (U.S.C.). However, the skills and target knowledge for effective NetA are best developed and honed during peacetime intelligence or network warfare support (NS) operations. Intelligence forces in the national intelligence community derive authority to conduct network exploitation and many NS [national security] operations from laws contained in U.S.C. Title 50. For this reason, ‘dual-purpose’ military forces are funded and controlled by organizations that derive authority under laws contained in both Title 10 and Title 50. The greatest benefit of these ‘dual-purpose’ forces is their authority to operate under laws contained in Title 50, and so produce actionable intelligence products while exercising the skills needed for NetA. These forces are the preferred means by which the Air Force can organize, train, and equip mission-ready NetA forces.” See http://www.herbb.hanscom.af.mil/tbbs/R1528/AF_Doctrine_Doc_2_5_Jan_11_2005.pdf.

tected than is the average system on the Internet. Nevertheless, as press reports in recent years make clear, such measures do not guarantee that their systems are immune to the hostile actions of outsiders.⁶⁷

As for gathering the intelligence needed to penetrate an adversary computer or network for cyberexploitation, this process is essentially identical to that for cyberattack. The reason is that cyberexploitation and cyberattack make use of the same kinds of access paths to their targets, and take advantage of the same vulnerabilities to deliver their payloads. In the event that an adversary detects these intelligence-gathering attempts, there is no way at all to determine their ultimate intent.

2.6.4.3 Rules of Engagement and Command and Control

Rules of engagement for cyberexploitation specify what adversary systems or networks may be probed or penetrated to obtain information. A particularly interesting question arises when a possible target of opportunity becomes known in the course of an ongoing cyberexploitation. For example, in the course of exploring one adversary network (Network A), the exploiter may come across a gateway to another, previously unknown network (Network B). Depending on the nature of Network B, the rules of engagement specified for Network A may be entirely inadequate (as might be the case if Network A were a military command and control network and Network B were a network of the adversary's national command authority). Rules of engagement for cyberexploitation must thus provide guidance in such situations.

In at least one way, command and control for cyberexploitation is more complex than for cyberattack because of the mandatory requirement of report-back—a cyberexploitation that does not return information to its controller is useless. By contrast, it may be desirable for a cyberattack agent or weapon to report to its controller on the outcome of any given attack event, but its primary mission can still be accomplished even if it is unable to do so.

Report-back also introduces another opportunity for the adversary to discover the presence of an exploiting payload, and thus the exploiter must be very careful in how report-back is arranged.

⁶⁷ For example, the Slammer worm attack reportedly resulted in a severe degradation of the Bank of America's ATM network in January 2003. See Aaron Davis, "Computer Worm Snarls Web: Electronic Attack Also Affects Phone Service, BOFA's ATM Network," *San Jose Mercury News*, January 26, 2003, available at <http://www.bayarea.com/mld/mercurynews/5034748.htm+atm+slammer+virus&hl=en>.

2.6.4.4 Effectiveness Assessment

The cyberexploitation analog to damage assessment for cyberattack might be termed effectiveness assessment. If a cyberexploitation does not report back to its controller, it has failed. But even if it does report back, it may not have succeeded. For cyberexploitation, the danger is that it has been discovered and that somehow the adversary has provided false or misleading information that is then reported back. Alternatively, the adversary may have compromised the report-back channel itself and inserted its own message that is mistaken for an authentic report-back message. (In a worst-case scenario, the adversary may use the report-back channel as a vehicle for conducting its own cyberattack or cyberexploitation against the controller.)

These scenarios for misdirection are not unique to cyberexploitation, of course—they are possible in ordinary espionage attempts as well. But because it is likely to be difficult for an automated agent to distinguish between being present on a “real” target versus being present on a “decoy” target, concerns about misdirection in a cyberexploitation context are all too real.

2.6.4.5 Tradeoffs Between Cyberattack and Cyberexploitation

In contemplating what to do about an adversary computer or network, decision makers have essentially two options—render it unavailable for serving adversary purposes or exploit it to gather useful information. In many cases, these two options are mutually exclusive—destroying it makes it impossible to exploit it. In some cases, destroying it may also reveal to the adversary some vulnerability or access path previously unknown to him, and thus compromise friendly sources and methods.

These tradeoffs are no less present in cyberattack and cyberexploitation. But in some ways, the tradeoffs may be easier to manage. For example, because a given instrument for cyberexploitation can be designed with cyberattack capabilities, the transition between exploitation and attack may be operationally simpler. Also, a cyberattack may be designed to corrupt or degrade a system slowly—and exploitation is possible as long as the adversary does not notice the corruption.

2.7 HISTORICAL PRECEDENTS AND LESSONS

To provide a sense of what might be possible through cyberattack and cyberexploitation, it is useful to consider some of the ways in which criminals have used them. A number of such cases are described in Appen-

dix C, and some of the lessons derived from considering these cases are provided below.

- Attacks can have multiple phases, as illustrated in several of the cases in Appendix C, that last over a relatively long period of time (over a year, in many cases.) This is especially true of DDOS attacks, where attackers must first take control of thousands and thousands of computers by installing their malicious software on them, causing them to join into mass command and control (e.g., join a botnet in IRC channels.) The same bots that are used for DDOS are also used for recruiting new bots through direct attack, sending copies of the malware to addressees in the victimized computer's address book. The less visible or "noisy" the activity, the longer the multiphase attack can last before being detected and mitigated.

- Attacks can also have multiple foci. In the Invita case (Appendix C), there was a primary focus on trying to locate credit card data to perpetrate fraud, but the attackers also used extortion to obtain financial gain. In some of the botnet cases, the botnets would be used for extortion or click-fraud. The Stakkato case was multitarget, but this was primarily a by-product of following login trust relationships between systems and sites.

- The same tactics used to compromise one host can be extended to compromise 1,000 hosts, given enough resources to repeat the same steps over and over, assuming the attacked systems are part of the same system monoculture all running the same targeted software (such as the same operating system). Automating these steps makes the job even easier, which can readily be done. (Anything that a user can do by typing at a keyboard can be turned into a scripted action. This is how the Invita attackers managed the creation and use of e-mail and online bidding accounts.)

A corollary is the notion that an indirect attack can be as successful as a direct attack, given the resources necessary to work through the entire set of login relationships between systems. For example, one can attempt to get access to another person's account by attacking that target's laptop or desktop system. This may fail, because the target may secure its personal computers very well. But the target may depend on someone else for system administration of its mail spool and home directory on a shared server. The attacker can thus go after a colleague's, a fellow employee's, or the service provider's computer and compromise it, and then use that access to go after an administrator's password on the file server holding the target's account.

The best case (from an attacker's standpoint) is when the same vul-

nerability exists at all levels within large interconnected systems, where “redundant” resources can be compromised, resulting in cascading effects.⁶⁸ This situation could allow an adversary to very quickly commandeer a large and diverse population of systems, as has been witnessed in various worm outbreaks over the past few years.

- The theft of credentials, either for login authentication or executing financial transactions, is a popular and successful avenue of attack. All that is necessary is either to direct a user to pass his or her keystrokes through a program under control of the attack (e.g., as in “phishing” attacks), or to get administrative control of either clients or servers and install software that logs keystrokes.

- Highly targeted attacks against specific companies are possible, as was seen in the Israeli industrial espionage case, as well as a variant of the BugBear trojan in 2003 that specifically targeted the domains of more than 1,000 specific banks in several countries.⁶⁹ Discovery and taking advantage of implicit business trust relationships between sites are also possible, as was seen in the Stakkato case. An attacker need only start with the most basic information that can be obtained about a company through open sources (e.g., press releases, organizational descriptions, phone directories, and other data made public through websites and news stories). She then uses this information to perform social engineering attacks, a pretext designed to trick users into giving out their passwords so that she can gain access to computers inside an organization’s network. Once in control of internal hosts, she effectively has insider access and can leverage that access to do more sensitive intelligence gathering on the target. She can learn business relationships, details about active projects and schedules, and anything necessary to fool anyone in the company into opening e-mail attachments or performing other acts that result in compromise of computer systems. (This is basic intelligence collection and analysis.) Control of internal hosts can also be used to direct attacks—behind the firewall and intrusion detection systems or intrusion prevention systems—against other internal hosts.

⁶⁸ See, for example, Daniel E. Geer, “Measuring Security,” 2006, pp. 170-178, available at <http://geer.tinho.net/measuringsecurity.tutorialv2.pdf>.

⁶⁹ F-Secure, “F-Secure Virus Descriptions: Bugbear.B,” 2003, available at http://www.f-secure.com/v-descs/bugbear_b.shtml.

Part II

Mission and Institutional Perspectives

Part II contains three different mission perspectives. Chapter 3 addresses military perspectives on cyberattack, largely from the point of view of the Department of Defense. Chapter 4 addresses intelligence perspectives on cyberattack (for conducting covert action) and cyberexploitation (for obtaining information). Chapter 5 addresses federal law enforcement perspectives on cyberattack and cyberexploitation and considers how the private sector might view cyberattack as an element of its defensive posture. These chapters depict in outline form how various institutions both inside and outside government have conceptualized or may in the future conceptualize the use of cyberattack and cyberexploitation technologies.

Regrettably, the picture that emerges from these chapters is fragmented and incomplete—largely because national policy with respect to cyberattack is fragmented and incomplete. The secrecy that surrounds policy in this area has further worsened the coherence of the overall picture. On the other hand, it is often true that with a new and easily available technology (the technology of cyber offensive actions), interests among a variety of different institutional actors in using this technology have arisen from the bottom up—from those with operational missions. In the early stages of technology adoption, some actors consider how the technology of cyber offensive actions might support or better enable the performance of their traditional missions—and others ignore it. The bottom-up nature of technology adoption in such cases inevitably leads to adoptions at different rates and for relatively parochial purposes, and

so the fragmentation of policy and organization today is not entirely surprising.

Nevertheless, the committee believes there is value in setting forth a notional view of how these institutions might conceptualize the uses of cyberattack, the associated decision-making structures, and the infrastructure needed to support the use of cyberattack as an instrument in their toolkits. If nothing else, the availability of a notional view provides a framework against which to react and within which to pose questions about what might be missing. These comments should be kept in mind as these chapters are read.

Part II also contains Chapter 6, a description of decision-making and oversight mechanisms in both the executive and the legislative branches that are relevant for cyberattack. Considering these mechanisms from a top-down perspective is intended to provide some points of reference that can help to identify what is missing from the picture painted by Chapters 3, 4, and 5.

A Military Perspective on Cyberattack

3.1 U.S. MILITARY DOCTRINE AND CYBERATTACK

The most current statement of U.S. military doctrine regarding cyberattack identifies computer network attack (an aspect of what this report calls cyberattack) as an element of computer network operations (CNO), the other two of which are computer network defense (CND) and related computer network exploitation (CNE) enabling operations. Computer network attack (CNA) refers to actions taken through the use of computer networks to disrupt, deny, degrade, or destroy information resident in computers and computer networks, or the computers and networks themselves. CND refers to actions taken through the use of computer networks to protect, monitor, analyze, detect, and respond to unauthorized activity against or within DOD information systems and computer networks. CNE (computer network exploitation, an aspect of what this report calls cyberexploitation) refers to operations conducted through the use of computer networks to gather data from target or adversary automated information systems or networks, and the term “related CNE enabling operations” refers to operations undertaken to gather intelligence information for carrying out CNO or CND operations.

Current doctrine (Joint Publication 3-13, *Joint Doctrine on Information Operations*) notes that all of these capabilities can be used for both offensive and defensive purposes. For example, under this rubric, a computer network attack might be used for a defensive purpose, such as the neutralization of a cyberthreat to a DOD computer or network.

At the date of this writing, an unclassified and authoritative state-

ment of current joint doctrine for the use of computer network attack is unavailable, and it is fair to say that current doctrine on this matter is still evolving. However, in testimony to the House Armed Services Committee on March 21, 2007, General James E. Cartwright, Commander of the United States Strategic Command, said that “cyberspace has emerged as a warfighting domain not unlike land, sea, and air, and we are engaged in a less visible, but nonetheless critical battle against sophisticated cyberspace attacks.” He pointed out the importance of deterring adversaries and assuring U.S. freedom of action in cyberspace, and argued that “fundamental to this approach is the integration of cyberspace capabilities across the full range of military operations.” He then observed that “to date, our time and resources have focused more on network defenses to include firewalls, antivirus protection, and vulnerability scanning. [But] while generally effective against unsophisticated hackers, these measures are marginally effective against sophisticated adversaries.”

Following this observation, he then stated:

History teaches us that a purely defensive posture poses significant risks; the “Maginot Line” model of terminal defense will ultimately fail without a more aggressive offshore strategy, one that more effectively layers and integrates our cyber capabilities. If we apply the principles of warfare to the cyber domain, as we do to sea, air, and land, we realize the defense of the nation is better served by capabilities enabling us to take the fight to our adversaries, when necessary to deter actions detrimental to our interests.

A number of other DOD and service statements and publications have added texture to the perspective articulated by General Cartwright. The 2006 *National Military Strategy for Cyberspace Operations* (redacted copy available online¹) says that “as a war-fighting domain . . . cyberspace favors the offense . . . an opportunity to gain and maintain the initiative.” It further defines cyberspace as a domain “characterized by the use of electronics and the electromagnetic spectrum to store, modify, and exchange data via networked systems and associated physical infrastructures.”

Prevailing military doctrine calls for the U.S. dominance of domains of warfare, traditionally including land, sea, air, and space, and now including cyberspace. Dominance in a domain means that the U.S. military should have freedom of access to and use of the domain, and should be able to deny access to and use of that domain to an adversary—and dominance requires that the United States play both offense and defense. Furthermore, if cyberspace is like any other warfighting domain, the fundamental concepts of warfare must apply to the cyberspace domain.

¹ See <http://www.dod.mil/pubs/foi/ojcs/07-F-2105doc1.pdf>.

An example of how such thinking regarding cyberspace-as-domain can play out was described to the committee in a briefing from the Air Force Cyberspace Task Force. In the CTF view, the United States should be provided with “offensive capabilities and deliberate target sets.” In addition, the briefing argued that “cyber favors the offensive” and that under this rubric fell several different missions, including strategic attack directly at enemy centers of gravity, suppression of enemy cyberdefenses, offensive counter-cyber, defensive counter-cyber, and interdiction. Consistent with Secretary of the Air Force Michael W. Wynne’s statement that “all aspects of air war will have some equivalent role in cyber war,”² these missions have rather close analogs to traditional Air Force missions—strategic bombing attack against enemy centers of gravity, suppression of enemy air defenses to facilitate airspace penetration of enemy borders, offensive counter-air (destroying enemy aircraft on the ground), defensive counter-air (defending friendly territory from enemy aircraft in the air), and interdiction (attack of enemy assets far behind the battlefield).³ (Whether this particular view of cyberspace as a domain of military conflict will ultimately be adopted throughout the Department of Defense is not clear at this time.)

The doctrinal perspective that cyberspace is another warfighting domain has other implications as well. For example, operations in cyberspace need to be synchronized and coordinated with other operations, just as land and air operations, for example, must be synchronized and coordinated. In other words, during overt or open military conflict, it is highly likely that information operations—including cyberattacks if militarily appropriate—will not be the only kind of military operations being executed. Examples of coordination issues are described in Box 3.1.

The doctrinal perspective further implies that cyberweapons should be regarded as no different from any other kind of weapon available to U.S. forces. That is, their use should be initiated on the basis of their suitability for conducting the attacks in question, and should not require any extraordinary analysis or authority to which the non-cyberspace military is not already accustomed. Thus, in determining the best way to attack a target, cyberweapons simply provide the operational planner with another option, in addition to the air-delivered laser-guided bomb and the Special Operations force with demolition charges.

Similar considerations apply from a legal perspective. For example,

² Michael W. Wynne, “Flying and Fighting in Cyberspace,” *Air & Space Power Journal*, Spring 2007, available at <http://www.airpower.maxwell.af.mil/airchronicles/apj/apj07/spr07/wynnespr07.html>.

³ Indeed, Lt. Gen. Bill Donahue (USAF, ret.) argued in a briefing to the committee that one could almost literally do a global search and replace that would replace “Air” with “Cyberspace” in Air Force warfighting doctrine.

BOX 3.1 Possible Coordination Issues for Cyberattack

Cross-domain coordination requires that the effects of a cyberattack on the physical world (both direct and consequential) and the timing of those effects should be known with enough certainty that their possible use can be taken into account in operational planning. Some issues include the following:

- *Coordination with other military operations.* Planners might choose to attack a given target using both a cyberweapon and a kinetic weapon. Redundancy in an attack, especially using different modes of attack that might exploit different vulnerabilities, is often desirable from a planning perspective. On the other hand, problems may result if the damage assessment from one operation is not available to those planning the other operation (e.g., as the result of stovepiping within executing agents).

- *Coordination between cyber operations for attack and for defense.* A computer network attack launched by the U.S. military may stimulate a counter-response from an adversary that could affect U.S. computers and networks, which may—or may not—be under military control. For example, a cyberattack that is conducted against a target in a given geographic command (e.g., PACOM) by the U.S. Strategic Command may stimulate action that has an impact on the regional networks used by that geographic command. A cyberattack launched by the United States may also stimulate adversary action that would have an impact on private sector network use and potentially disrupt important civilian activities—suggesting that cyberattacks by the U.S. military may have defensive implications.

- *Coordination between cyberattack and cyberexploitation.* Unless attack and exploitation are coordinated, it is possible to imagine scenarios in which a cyberattack to plant false information in an adversary's database results in the cyberexploitation extracting that false information and using it as though it were real and valid. And, of course, there is the classic conflict about whether it is more desirable to shut down an adversary's communication channel (an attack operation) or to listen to it (an exploitation operation).

all military operations are subject to certain limitations mandated by the law of armed conflict regarding differentiation of targets, military necessity, limiting collateral damage, and so on. Of course, targets in cyberspace are different from targets on the ground, so the facts relevant to any given operation may be different in the former case than in the latter, but the analytical process remains the same. Thus, if it was legitimate to attack a target with kinetic weapons, it remains legitimate under the laws of armed conflict to attack it with cyberweapons. These considerations are addressed at length in Chapter 7.

In short, according to this perspective, conflict in cyberspace should be treated like conflict in a physical domain, the same rules and policies should apply, and the only differences are operational.

3.2 DEPARTMENT OF DEFENSE ORGANIZATION FOR CYBERATTACK

The U.S. Strategic Command (STRATCOM) plays a key role in DOD cyber operations. STRATCOM is composed of eight functional components, including five Joint Functional Component Commands (JFCCs).⁴ Each JFCC is responsible for focusing on a specific operational area—one of those operational areas involves offensive network warfare (NW) and defensive network operations (NetOps).⁵

Offensive network warfare is the responsibility of the Joint Functional Component Command for Network Warfare (JFCC-NW). The commander of the JFCC-NW is also the director of the National Security Agency (NSA) and is “responsible for deliberate planning of network warfare, which includes coordinated planning of offensive network attack.”⁶ JFCC-NW was established in January 2005.⁷ Network warfare as used in the context of JFCC-NW means “the employment of Computer Network Operations (CNO) with the intent of denying adversaries the effective use of their computers, information systems, and networks, while ensuring the effective use of our own computers, information systems, and networks.”⁸ These operations include computer network attack (CNA), computer network exploitation (CNE), and Computer Network Defense (CND). The JFCC-NW also supports the network warfare needs of Combatant Commands/Commanders (COCOMs).⁹

Defensive network operations are the responsibility of the Joint Task Force-Global Network Operations (JTF-GNO). The commander of JTF-GNO is also the director of the Defense Information Systems Agency and is responsible for operating and defending the DOD information infra-

⁴ The eight components are JFCC-Global Strike and Integration (JFCC-GSI), JFCC-Integrated Missile Defense (JFCC-IMD), JFCC-Intelligence, Surveillance and Reconnaissance (JFCC-ISR), JFCC-Space (JFCC-SPACE), Joint Information Operations Warfare Command (JIOWC), STRATCOM Center for Combating Weapons of Mass Destruction (SCC-WMD), and Joint Task Force-Global Network Operations (JTF-GNO). See http://www.stratcom.mil/organization-fnc_comp.html.

⁵ Lt. Gen. Keith B. Alexander, “Warfighting in Cyberspace,” *Joint Force Quarterly* 46(3):58-61, 2007.

⁶ Clay Wilson, “Information Operations and Cyberwar: Capabilities and Related Policy Issues,” U.S. Congressional Research Service (RL31787), updated September 14, 2006, p. 8.

⁷ JFCC-NW Implementation Directive, January 20, 2005. Cited in Keith B. Alexander, “Warfighting in Cyberspace,” *Joint Force Quarterly*, July 2007, available at <http://www.military.com/forums/0,15240,143898,00.html>.

⁸ USSTRATCOM Command Video, available at <http://www.stratcom.mil/Videos/transcripts/Command%20Video.txt>.

⁹ Joint Publication 3-13 (2006) states that STRATCOM has responsibility for “identifying desired characteristics and capabilities of CNA, conducting CNA in support of assigned missions, and integrating CNA capabilities in support of other combatant commanders.”

structure known as the Global Information Grid (GIG). The Joint Information Operations Warfare Command, responsible for assisting combatant commands with an integrated approach to information operations, coordinates network operations and network warfare with the JTF-GNO and the JFCC-NW.¹⁰ As of November 2008, the JTF-GNO is for the first time placed under the operational control of the JFCC-NW.¹¹

The JFCC-NW engages in a substantial amount of coordination with other entities. It coordinates its offensive activities directly with the defensive activities of the JTF-GNO. It “facilitates cooperative engagement with other national entities in computer network defense and network warfare as part of global information operations.”¹² Because the commander of the JFCC-NW is dual-hatted as the director of the National Security Agency (Box 3.2), the JFCC-NW can easily work with the intelligence community to provide intelligence support for computer network operations. In addition, coordination between cyberattack (a Title 10 function) and cyberexploitation (a Title 50 function) is more easily accomplished.

Lastly, Joint Publication 3-13 also notes that

CDRUSSTRATCOM’s specific authority and responsibility to coordinate IO [information operations, Box 3.3] across AOR and functional boundaries does not diminish the imperative for the other combatant commanders to coordinate, integrate, plan, execute, and employ IO. These efforts may be directed at achieving national or military objectives incorporated in TSCPs [Theater Security Cooperation Programs], shaping the operational environment for potential employment during periods of heightened tensions, or in support of specific military operations.

Two important points are embedded in this paragraph. First, STRATCOM is not necessarily the only command that can actually carry out information operations, including computer network attack. (In some cases, STRATCOM will be a supporting command that provides support to other regional or functional commands. In other cases, it will be the supported command, receiving support from other regional or functional commands.) Second, information operations, including computer network attack, may be used both in support of specific military operations *and* during periods of “heightened tensions,” that is, early use *before* overt conflict occurs.

¹⁰ Clay Wilson, “Information Operations and Cyberwar,” 2006.

¹¹ Memo of Robert Gates (Secretary of Defense) to DOD regarding Command and Control for Military Cyberspace Missions, November 12, 2008. Copy available from the NRC.

¹² U.S. Strategic Command website, <http://www.stratcom.mil/about-ch.html>.

BOX 3.2 The National Security Agency-Central Security Service

Often known simply as the National Security Agency, the organization is in fact a combat support agency of the DOD under the authority, direction, and control of the Secretary of Defense, and is responsible for centralized coordination, direction, and performance of highly specialized intelligence functions in support of U.S. government activities. It includes both the National Security Agency and the Central Security Service. The NSA carries out the responsibilities of the Secretary of Defense to serve as executive agency for U.S. government signals intelligence (SIGINT), communications security, computer security, and operations security training activities. The CSS is composed of the Service Cryptologic Elements of the four uniformed services that are responsible for conducting their Title 50 SIGINT mission, and provides the military Services a unified cryptologic organization within the DOD that assures proper control of the planning, programming, budgeting, and expenditure of resources for cryptologic activities. Service cryptologic elements also perform other missions in direct support of their respective Services related to information operations (including computer network operations), and in doing so, they operate under Title 10 authority.

The director of the National Security Agency (DIRNSA) serves as the director of both the National Security Agency and the Central Security Service and has both Title 10 and Title 50 responsibilities. As national executive agent for SIGINT, DIRNSA has operated with Title 50 authority and thus would be responsible for conducting cyberexploitations, which by definition are not supposed to damage, degrade, or disable adversary computer systems or networks. As the party responsible for DOD information assurance, DIRNSA has operated with Title 10 authority. Finally, in January 2005, the Joint Functional Component Command for Network Warfare (JFCC-NW) was established under the U.S. Strategic Command, and DIRNSA was designated as its commander. As such, DIRNSA operates with Title 10 authority for any offensive missions (including cyberattacks) undertaken by the JFCC-NW.

As this report is being written, these arrangements are in flux, as the DOD and the intelligence community are discussing the potential standup of a cyber combatant command.

3.3 RULES OF ENGAGEMENT

In general, the rules of engagement (ROEs) for military forces specify the circumstances under which they may take certain kinds of action. (The laws of armed conflict place additional constraints on the permissible actions of military forces.) For example, many military installations contain areas in which “the use of deadly force is authorized” to stop individuals from trespassing—guards of such areas are authorized (but not required) to use any means necessary to do so.

BOX 3.3 Information Operations and Related Capabilities

Computer network operations are themselves part of a larger complex designated as information operations (IO) by the Joint Chiefs of Staff. These other elements of information operations include:

- *Psychological operations (PSYOP)*, which include operations to convey selected truthful information and indicators to foreign audiences to influence their emotions, motives, objective reasoning, and ultimately, the behavior of their governments, organizations, groups, and individuals. The purpose of PSYOP is to induce or reinforce foreign attitudes and behavior favorable to the originator's objectives.
 - *Military deception*, which includes actions taken with the purpose of deliberately misleading adversary decision makers as to friendly military capabilities, intentions, and operations, thereby causing the adversary to take specific actions (or inactions) that will contribute to the accomplishment of the friendly forces' mission.
 - *Operations security (OPSEC)*, which is a process of identifying critical information and subsequently analyzing friendly actions and other activities to identify what friendly information is necessary for the adversary to have sufficiently accurate knowledge of friendly forces and intentions; deny adversary decision makers critical information about friendly forces and intentions; and cause adversary decision makers to misjudge the relevance of known critical friendly information because other information about friendly forces and intentions remains secure. OPSEC seeks to deny real information to an adversary and prevent correct deduction of friendly plans.
 - *Electronic warfare (EW)* refers to any military action involving the use of electromagnetic (EM) and directed energy to control the EM spectrum or to attack the adversary. EW includes electronic attack (EM energy, directed energy, or antiradiation weapons to attack personnel, facilities, or equipment with the intent of degrading, neutralizing, or destroying adversary combat capability), electronic protection (which ensures the friendly use of the EM spectrum), and electronic warfare support (ES, which searches for, intercepts, identifies, and locates or localizes sources of intentional and unintentional radiated EM energy for the purpose of immediate threat recognition, targeting, planning, and conduct of future operations). ES data can be used to produce SIGINT, provide targeting for electronic or other forms of attack, and produce measurement and signature intelligence (MASINT). SIGINT and MASINT can also provide battle damage assessment (BDA) and feedback on the effectiveness of the overall operational plan.

In addition, a number of other capabilities support information operations in the DOD context, such as information assurance (IA), physical security, physical attack, and counterintelligence. Capabilities related to IO include public affairs (PA), civil-military operations (CMO), and defense support to public diplomacy. The Joint Chiefs of Staff note that these capabilities can also make significant contributions to IO but that their primary purpose and the rules under which they operate must not be compromised by IO.

Some of the issues relevant to formulating ROEs for cyberattack might include:

- When to execute a cyberattack—what are the circumstances under which a cyberattack might be authorized?
- Scope of a cyberattack—what are the entities that may be targeted?
- Duration of the cyberattack—how long should a cyberattack last?
- Notifications—who must be informed if a cyberattack is conducted?
- Authority for exceptions—what level of authority is needed to grant an exception for standing ROEs?

To illustrate, consider the standing rules of engagement promulgated by the Joint Chiefs of Staff, which state that “a [U.S.] commander has the authority and obligation to use all necessary means available and to take all appropriate [i.e., necessary and proportional] actions to defend that commander’s unit and other U.S. forces in the vicinity from a hostile act or *demonstration of hostile intent* [emphasis added]”¹³ where “hostile intent” is understood to mean that another party has taken some action that reasonably indicates a potential for more or less immediate attack.

Applying this rule to the cyber domain raises the question of actions that constitute a demonstration of hostile intent. For example, do non-destructive adversary probes of important military U.S. computer systems and networks (or even systems and networks associated with U.S. critical infrastructure) constitute demonstrations of hostile intent? If so, do such actions justify actions beyond the taking of additional passive defense measures? Would a commander be permitted to conduct probes on adversary networks from which these probes were emanating? To conduct a responsive cyberattack to neutralize the probes?

On this specific topic, Rear Admiral Betsy Hight of the Joint Task Force on Global Network Operations testified to the committee that the commander of the U.S. Strategic Command has operational authority to conduct cyber operations that are defensive in purpose against systems outside the DOD networks. The action taken in the operation may have an offensive character—that is, it may seek to damage or disrupt a system that is adversely affecting a DOD asset. Self-defense is generally limited in scope to addressing or mitigating the immediate hostile act, and is a last resort. The frequency with which the U.S. Strategic Command has actually acted under this asserted authority, if at all, is unknown.

¹³ Joint Chiefs of Staff, Chairman of the Joint Chiefs of Staff Instruction, CJCSI 3121.01A, January 15, 2000, Standing Rules of Engagement for US Forces, available at http://www.fas.org/man/dod-101/dod/docs/cjcs_sroe.pdf.

CND response actions (RAs) are a specific subset of self-defense and are likewise constrained to a measured response used as a last resort. CND RAs can be used only in response to a network event that creates a threshold impact. Additional limitations constrain the scope, duration, and impact of the CND RA. Moreover, CND RAs, like all self-defense, is a tactical activity, characterized as such because it is used in response to a specific hostile action and is designed to address and mitigate that action, and only that action. Offensive actions are not so limited. Both offensive and defensive actions must follow the law of war limitations with regard to differentiation of targets, military necessity, and limiting collateral damage, but defensive actions tend to be more limited in scope.

Such self-defense operations would be designated as a CND response action, authority for which is described, constrained, and granted through standing rules of engagement established by the National Command Authority and flowing, through the secretary of defense, from the President's authority as commander-in-chief. Standing rules of engagement generally describe the authority commanders have to defend their personnel and designated property. According to Admiral Hight's testimony to the committee, the rules of engagement for CND response actions also specify that they are not authorized unless the hostile action has an impact on the ability of a combatant commander to carry out a mission or an ongoing military operation, and in particular that hostile actions that result only in inconvenience or that appear directed at intelligence gathering do not rise to this threshold.

An example of a legitimate target for a CND response action would be a botnet controller that is directing an attack on DOD assets in cyberspace. Thus, if bots are active in DOD networks, and if through DOD mission partners the controller of those bots can be identified in cyberspace, and if the botnet attack is compromising the DOD network's ability to carry out its mission operationally, a CND response action—involving cyberattack—can be directed against the controller under these standing rules of engagement.

As for geographic scope, a hostile cyber act may emanate from anywhere in cyberspace. Accordingly, the impact of CND response actions directed against that source could also occur anywhere in cyberspace. The ease with which actors can use and misuse U.S.-based cyber assets for malicious purposes increases the probability that future CND response actions might impact that space. For this reason, the JTF-GNO maintains relationships with law enforcement, other federal entities, and Internet service providers. This ensures that if some other national asset, or the commercial sector, can mitigate malicious cyber activity against the DOD, those assets are used before resorting to CND response actions.

The final point about this particular example is that from the DOD

perspective, the cessation of a hostile action may be more important than the attribution of the action to a particular actor. Accordingly, under the stated policy, the DOD may be willing to take many steps to ensure that the hostile action ceases, even if those actions have ramifications beyond U.S. borders.

3.4 SOME HISTORICAL PERSPECTIVE

Because the number of confirmed and unclassified instances of cyberattack launched by governments, friendly or hostile, is vanishingly small, it is hard to cite actual experience as a basis for understanding the effects of cyberattack. But a number of other incidents can provide some insight. Although the events described are not cyberattacks themselves, the affected entities involved are the kinds of targets that proponents of cyberattack weapons often discuss when advancing the case for the value of such weapons. The operational effects are the kinds of effects that cyberattacks might seek to cause.

- In December 2006, a major fiber optic cable providing some 50 percent of Iran's digital communications and Internet connections was damaged in Iran's territorial waters in the Persian Gulf. A month later, 80 percent of the damaged capability had been restored.
- In late December 2006, an earthquake off the shores of Taiwan damaged or destroyed eight fiber optic lines that connected Taiwan to other nations in the Pacific. There was some disruption to Internet and phone for about 2 days, and Internet connections were slow in Taiwan, Hong Kong, Japan, China, Singapore, and South Korea. However, although the cables were not repaired for almost 3 weeks, workaround restored most services quickly.
- In February 2007, Mexico's largest cell phone company experienced a "crash" that left 40 million cell phone users without service for most of a day.
- In May 1999, the United States targeted the Belgrade electric power system as part of the Kosovo conflict, using carbon fibers to short generators. In all, four strikes were conducted against the power system, but in each case, power generation was restored within a few days to a substantial fraction of what it was prior to the strike.

Perhaps the most important feature of these incidents is the fact that their effects were relatively transitory, largely because the parties affected found workarounds that enabled them to compensate for the immediate loss of capability. If these incidents had been caused deliberately, it is likely that repeated attacks would have been necessary to ensure that

the reduction of capability persisted over time. Moreover, these incidents were, by themselves, of little strategic significance, though if they had been timed to coincide with some kinetic military operation, they might well have had a significant impact.

At the same time, these observations do not account for possible impact on the psychological state of mind of relevant decision makers. The same major outage of service may result from a natural disaster, by a deliberate overt action, or by a deliberate and well-concealed action—but decision makers are likely to care about the specific cause of such an event. An outage caused by deliberate action is fundamentally different from the “same” outage caused by natural disaster, because the first carries with it the implicit threat of happening again when an adversary wants it to happen again. A well-concealed action attributed to a specific party could be argued to have an even greater impact on a decision maker, since he might well be hard-pressed to do anything about it.

The 2007 cyberattack on Estonia yielded similar lessons (Box 3.4). The attack had a variety of short-term consequences for Estonia, including the inability of Estonians to access online banking services and government services, and for individuals outside Estonia to access the Estonian web for a while. Less measurable impacts such as confusion and miscommunication were also noted.¹⁴ Although these impacts led press reports to suggest that the conflict was variously the first war in cyberspace,¹⁵ Web War I,¹⁶ and “the first real example of nation-states flexing their cyber-warfare capabilities,”¹⁷ no critical infrastructure was targeted in the attacks, most sites were restored to service quickly, and the primary operational result of the attack was inconvenience.

This is not to say that the attack was inconsequential. The incident did serve as a “wake-up call” for many other nations to inquire how they should respond to similar situations that might arise in the future. During the attack, NATO provided experts in Internet warfare to assist in the investigation and defense.¹⁸ Furthermore, in the aftermath of the attacks, Estonia has proposed that NATO create a Cooperative Cyber Defense Center of Excellence to improve NATO members’ ability to cooperate

¹⁴ Jaak Aaviksoo, Minister of Defense of Estonia, presentation to Centre of Strategic and International Studies, November 28, 2007, p. 3. of transcript, available at http://www.csis.org/component?option=com_csis_press/task/view/id,3525/.

¹⁵ Mark Landler and John Markoff, “In Estonia, What May Be the First War in Cyberspace,” *International Herald Tribune*, May 28, 2007.

¹⁶ Joshua Davis, “Hackers Take Down the Most Wired Country in Europe,” *Wired*, Issue 15.09, August 21, 2007.

¹⁷ McAfee Corp., “Cybercrime: The Next Wave,” McAfee Virtual Criminology Report, 2007, p. 9.

¹⁸ *Economist*, “A Cyber-riot,” May 10, 2007.

BOX 3.4 The Cyberattacks on Estonia and Georgia

Estonia

On April 27, 2007, a series of distributed denial of service (DDOS) attacks began on a range of Estonian government websites, media sites, and online banking services.¹ Attacks were largely conducted using botnets to create network traffic. The duration and intensity of attacks varied across the websites attacked. According to data collected by Arbor Networks, the attacks were primarily Internet Control Message Protocol (ICMP) floods with most lasting from 1 minute to 1 hour and a few lasting up to 10 hours. Most attacks had an intensity of 30 Mbps or less, though some measured up to 95 Mbps.² Some Estonian websites were also defaced by people claiming to be Russian hackers, and tools in the form of scripts to conduct attacks were offered on Russian hacker sites and chat rooms.³ Computers running those scripts became packet sources, also contributing to the attacks.⁴

The attacks followed the removal the previous night of a statue memorializing WWII Soviet war dead from the center of the Estonian capital of Tallinn. They continued off and on until mid-May after peaking on May 9th, the day Russia commemorates Victory in Europe.⁵ The attacks were started and stopped deliberately by the attackers rather than being shut down through defensive measures.⁶ The Estonian government was quick to claim links between those conducting the attacks and the Russian government.⁷ The Estonian minister of defense stated that the attacks were “unusually well-coordinated and required resources unavailable to common people.”⁸ He claimed this indicated involvement beyond the capabilities of outraged citizens, though he did not make any explicit claims about involvement by state actors. One expert in cyberterrorism was quoted as saying that the attacks bore the hallmarks of a “false flag” operation, used to test out defenses.⁹ Russian officials denied any involvement.¹⁰

Evidence of Russian involvement was circumstantial with no “smoking gun” found to indicate any connection between the Russian government and the conduct of the attacks.¹¹ Hillar Aareleid, chief security officer for Estonia’s version of the U.S. Computer Emergency Response Team, dismissed claims that a Russian government link could be proven.¹² The botnets were composed of compromised computers from the United States, Europe, Canada, Brazil, Vietnam, and other countries around the world. There was evidence of Russian nationalists promoting the attacks through blog posts with scripts and instructions for conducting DDOS attacks on Estonian websites.¹³ One script used in the attacks which sent ping floods to Estonian websites was shared extensively on Russian language boards.¹⁴ Some attackers in the earliest attacks were identified by their IP addresses as coming from Russia, including some from Russian state institutions.¹⁵ An Estonian news site stated that a member of Nashi, a Russian youth group tied to Russian President Putin, claimed that the group was behind the attacks, but there was no corroboration of this claim.¹⁶

Continued

BOX 3.4 Continued

Georgia

In August 2008, a military conflict involving land, air, and sea forces of Georgia and Russia occurred in South Ossetia and Abkhazia, provinces under the nominal control of Georgia. Russian military action in this conflict was immediately preceded by a number of cyberattacks against a variety of websites of the Georgian government.¹⁷ These attacks defaced these websites and also made it very difficult for the Georgian government to put out its side of the story. Cyberattacks against certain Georgian government Web sites reportedly continued even after Russia declared a cease-fire.

In broad outline, the cyberattacks against Georgia were very similar to those against Estonia. As in the Estonian case, these attacks were not conclusively traced to the Russian government, and the Russian government denied involvement.¹⁸ Various analysts argue that they were controlled by the Russian Business Network,¹⁹ a business organization alleged to have criminal ties, and even private Russian citizens.²⁰

The primary significance of the cyberattacks on Georgia is in their appearing to be the first instance of simultaneous actions involving cyberattack and kinetic attack, rather than in any of the particulars of the cyberattacks themselves.

¹ *Economist*, "A Cyber-riot," May 10, 2007; Jaak Aaviksoo, Minister of Defense of Estonia presentation to Centre of Strategic and International Studies, November 28, 2007.

² The most detailed measurements on the attacks are from Arbor Networks; Jose Nazario, "Estonian DDoS Attacks—A Summary to Date," May 17, 2007, available at <http://asert.arbornetworks.com/2007/05/estonian-ddos-attacks-a-summary-to-date/>. Those measurements also show a small percentage of the TCP SYN attacks.

³ Some examples are available from the F-Secure weblog at <http://www.f-secure.com/weblog/archives/archive-052007.html#00001188>. See also Miska Rantanen, "Virtual Harassment, But for Real," *Helsingin Sanomat*, May 6, 2007, available at <http://www.hs.fi/english/article/Virtual+harassment+but+for+real-/1135227099868>.

⁴ *Heise Security*, "Estonian DDoS—A Final Analysis," May 31, 2007, available at <http://www.heise-security.co.uk/news/90461>. This article quotes Jose Nazario from Arbor Networks. See also the Arbor Networks measurements cited previously.

⁵ Michael Lesk, "The New Front Line: Estonia under Cyberassault," *IEEE Security & Privacy* 5(4):76-79, 2007.

in operational situations and to develop a doctrine for responding to cyberattacks.¹⁹ Data on and analysis of the attacks have been provided to NATO members to inform efforts aimed at better defending against such

¹⁹ Jaak Aaviksoo, Minister of Defense of Estonia, presentation to Centre of Strategic and International Studies, November 28, 2007, p. 7. of transcript, available at http://www.csis.org/component?option=com_csis_press/task/view/id,3525/.

⁶ McAfee Corporation, "Cybercrime: The Next Wave," *McAfee Virtual Criminology Report*, 2007, p. 11.

⁷ Maria Danilova, "Anti-Estonia Protests Escalate in Moscow," *Washington Post*, May 2, 2007, available at http://www.washingtonpost.com/wp-dyn/content/article/2007/05/02/AR2007050200671_2.html. The article quotes both the Estonian president and ambassador to Russia as claiming Kremlin involvement.

⁸ Jaak Aaviksoo, minister of defense of Estonia, in a presentation at the Centre of Strategic and International Studies, November 28, 2007, p. 2 of transcript, available at http://www.csis.org/component/option,com_csis_press/task,view/id,3525/.

⁹ McAfee Corp., op. cit., p. 9. The report quotes Yael Shahar, International Institute for Counter-Terrorism, Israel.

¹⁰ McAfee Corp., op. cit., p. 7.

¹¹ E-mail from Jose Nazario of Arbor Networks, July 5, 2007. See also *Heise Security*, "Estonian DDoS—A Final Analysis," May 31, 2007, available at <http://www.heise-security.co.uk/news/90461>.

¹² Jeremy Kirk, "Estonia Recovers from Massive DDoS Attack," *Computerworld Security*, May 17, 2007, available at <http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9019725>.

¹³ Jeremy Kirk, "Russian Gov't Not Behind Estonia DDOS Attacks: Analysis Throws Doubt on Whether a Single Agency Alone Was Involved," *InfoWorld*, June 1, 2007, available at http://www.infoworld.com/article/07/06/01/Russia-not-behind-Estonia-DDOS-attacks_1.html.

¹⁴ *Heise Security*, op. cit.

¹⁵ Ian Traynor, "Russia Accused of Unleashing Cyberwar to Disable Estonia," *Guardian*, May 17, 2007, available at <http://www.guardian.co.uk/russia/article/0,,2081438,00.html>.

¹⁶ Cory Doctorow in a June 2, 2007, blog entry on Boing Boing (<http://www.boingboing.net/2007/06/02/estonia-didnt-suffer.html>) cited an Estonian news article from Postimees.ee posted on May 29, 2007, available at <http://www.postimees.ee/290507/esileht/siseuudised/263405.php>. See Owen Matthews and Anna Nemtsova, "Putin's Powerful Youth Guard," *Newsweek*, May 28, 2007, for a description of Nashi and its link to President Putin and the Russian government.

¹⁷ "Georgia Accuses Russia of Coordinated Cyberattack," *CNET News*, August 11, 2008, available at http://news.cnet.com/8301-1009_3-10014150-83.html.

¹⁸ John Markoff, "Before the Gunfire, Cyberattacks," *New York Times*, August 13, 2008, available at <http://www.nytimes.com/2008/08/13/technology/13cyber.html?fta=y>.

¹⁹ Gregg Keizer, "Cyberattacks Knock Out Georgia's Internet Presence," *Computerworld*, August 11, 2008, available at <http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9112201>.

²⁰ Byron Acohido, "Some Russian PCs Used to Cyberattack Georgia," *USA Today*, August 17, 2008, available at http://www.usatoday.com/tech/news/computersecurity/hacking/2008-08-17-russia-georgia-war-hackers_N.htm.

attacks.²⁰ From a legal and policy standpoint, the attack raised questions about whether such an attack constituted an armed attack in the sense intended by the UN Charter and whether cyberattacks against a member nation ought to be included in the provisions of Article V of the North

²⁰ McAfee Corporation, "Cybercrime: The Next Wave," *McAfee Virtual Criminology Report*, 2007, p. 11.

Atlantic Treaty which provides for collective self-defense if any member is attacked.²¹

A second set of issues appears to have emerged from the U.S. experience more generally with information operations in the Kosovo conflict. Analysts and decision makers considered using information operations, including computer network attack, as part of an integrated campaign against certain targets in Kosovo. However, in practice, options such as computer network attack proved harder to use than expected, in part because of the difficulties in obtaining the necessary approvals and authorizations to use them. In some cases, the approval process took so long that the utility of the operation had passed, at least in part because the execution of a particular option had many unknowns about likely effects. In other cases, it would have been relatively straightforward for the adversary to counter the use of such an option. The summary assessment of a senior military officer regarding information operations in Operation Allied Force—"a big wind-up to an underhand throw."

Of course, the past may not be the best predictor of the future, especially when Allied forces are just starting to explore the possibilities and limitations of information operations, and in particular where decision-making processes have not yet fully accommodated the need to account for information operations. These observations are offered only to suggest that initial predictions of easy application are not likely to be realized.

The past decade has also seen a number of shifts in doctrinal perspective. For example, in 1998 the DOD publication JP3-13, *Joint Doctrine for Information Operations*, made reference to offensive and defensive information operations, as well as to "information warfare." The 2006 revision of JP3-13, *Information Operations*, discontinued the terms "offensive IO" and "defensive IO" but retained the recognition that information operations can be applied to achieve both offensive and defensive objectives, and it eliminated the term "information warfare" from joint IO doctrine. Furthermore, it defined five core capabilities for information operations (electronic warfare, computer network operations, psychological operations, operations security, and military deception) and their associated supporting and related capabilities. Lastly, it established the core IO capability of computer network operations, integrating computer network attack, computer network defense, and computer network exploitation under one umbrella.

²¹ Ian Traynor, "Russia Accused of Unleashing Cyberwar to Disable Estonia," *Guardian*, May 17, 2007, available at <http://www.guardian.co.uk/russia/article/0,,2081438,00.html>.

3.5 CYBERATTACK IN SUPPORT OF MILITARY OPERATIONS— SOME HYPOTHETICAL EXAMPLES

What are some of the applications of cyberattack? It is helpful to consider several broad categories separately. Cyberattack can support information operations within the information operations sphere and also other military operations. In addition, cyberattack can be applied to missions that are not traditionally within the military domain.

3.5.1 Cyberattack in Support of Defense, Exploitation, and Other Information Operations

As noted above, cyberattack can be used defensively to eliminate a threat to DOD systems or networks (an application of computer network defense). For example, the DOD might use a botnet to launch a DDOS counterattack to disable the computers from which a threat to DOD systems originates.²² In support of CNE, a cyberattack could be used to disable security software so that a cyberexploitation could insert monitoring software (e.g., key loggers) on adversary computers or networks.

Cyberattack can also be used to support other non-computer IOs. For example:

- *Psychological operations.* A cyberattack could be used to generate frequent e-mail messages or telephone calls to specific adversary decision makers. The frequency of such e-mail messages or phone calls could disrupt their work environments, making it difficult for them to work there. And the content of such e-mail messages could include threats such as “your building is going to be bombed in 30 minutes; it is a good idea if you leave” or “we know where your lover’s safe house is.”²³ Another PSYOP application might call for the launching of a small but very visible

²² The notion that the United States would actually do so—use a botnet in such a manner—is speculative, but such speculation has been seen from senior military lawyers, such as the staff judge advocate for the Air Force Intelligence, Surveillance and Reconnaissance Agency. See Charles W. Williamson III, “Carpet Bombing in Cyberspace: Why America Needs a Military Botnet,” *Armed Forces Journal International*, May 2008, available at <http://www.armedforcesjournal.com/2008/05/3375884>.

²³ *Air Force Doctrine Document 2-5* (issued by the Secretary of the Air Force, January 11, 2005) explicitly notes that “psychological operations can be performed using network attack [defined as employment of network-based capabilities to destroy, disrupt, corrupt, or usurp information resident in or transiting through networks] to target and disseminate selected information to target audiences.” See http://www.herbb.hanscom.af.mil/tbbs/R1528/AF_Doctrine_Doc_2_5_Jan_11_2005.pdf.

cyberattack and then announcing it to an adversary in order to undermine the adversary's confidence in its essential systems.²⁴

- *Operations security.* Cyberattacks could be used to target specific adversary sensor systems that are intended to report on information related to the location of friendly forces. For example, an adversary may have compromised a computer system on a DOD network that has access to information related to troop movements. An attack on that computer could render it inoperative, but it might be more useful to feed it incorrect information about troop movements knowing that such information might be highly trusted by the adversary.

- *Military deception.*²⁵ Cyberattacks could be used to gain access to an adversary computer system in its command and control structure. By assuming control of a computer used by a senior intelligence analyst, for example, bogus e-mail traffic could be sent to that analyst's customers. The contents of these e-mails could easily provide misinformation regarding the military capabilities, intentions, locations, and operations of friendly forces. Moreover, responding e-mails back to the analyst could be intercepted and appropriately modified before being displayed to the analyst.

- *Electronic warfare.* Cyberattacks could be used to disable an adversary's software-defined radios, thus preventing enemy wireless battlefield communications (which is often a goal of EW). In addition, EW could support cyberattacks. For example, to the extent that adversary computer systems are connected through wireless links, EW might be used to jam those links in order to disrupt the wireless network—that is, jamming would be a denial-of-service cyberattack against the network in question.

Cyberattack can also be used to support related missions, such as propaganda. Here is one possible example:

- Ruritania and Zendia are adversaries. Ruritania penetrates a Zendian GIS system focused on Armpitia, a Ruritarian ally, to alter maps and targeting databases. An Armpitian building containing a day-care center is marked as a munitions bunker, a historic cathedral as a troop barracks, and the embassy of a neutral nation as a branch of the ally's

²⁴ Defense Science Board, "Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software," U.S. Department of Defense, September 2007, p. 22.

²⁵ Air Force Doctrine Document 2-5 (issued by the secretary of the Air Force, January 11, 2005) explicitly notes that "network attack may support deception operations against an adversary by deleting or distorting information stored on, processed by, or transmitted by network devices." Available at http://www.herbb.hanscom.af.mil/tbbs/R1528/AF_Doctrine_Doc_2_5_Jan_11_2005.pdf.

ministry of defense. When Zendia launches an attack on Armpitia using cruise missiles, it destroys the embassy and the church, and kills dozens of children. CNN shows the evidence of the war crimes to the world. Public opinion swings against Zendia, war crime charges are filed at the Hague, and Zendian planners lose confidence in their standoff weapon systems.

Another example is the use of botnets to send spam e-mail carrying propaganda messages to an entire population. One related instance occurred in 2000, when a virus was used to spread information regarding a specific ethnically based incident or community in Sri Lanka.²⁶

3.5.2 Cyberattack in Support of Traditional Military Operations

Cyberattacks could also be used in connection with a variety of traditional military operations. Five illustrative examples are provided below:

- *Disruption of adversary command, control, and communications.* Such disruption could involve denial of service (so that those links are unusable) or spoofing or impersonation of legitimate authorities (so that information received by one party is not the information sent by the originating party). Tactical C2 networks and/or links between the adversary national command authority and forces in the field could be disrupted. Adversary planning (e.g., for military actions against U.S. forces) could be disrupted or altered clandestinely.

- *Suppression of adversary air defenses.* A networked air defense system that can pass data from forward-deployed sensors to air defense forces in the rear is much more effective than one without such coordination available. Disruption of such communications links can degrade the performance of the overall system considerably. It is also possible to imagine that long before any attack took place, an air defense radar delivered to an adversary might be clandestinely programmed to ignore certain radar signatures, namely those associated with airplanes friendly to the attacker, but only during certain times of day. From the adversary's perspective, the radar would appear to be working properly, as it would detect most airplanes most of the time. But the attacker would know the proper window to attack so that its airplanes would be ignored.

- *Degradation of adversary smart munitions and platforms (example 1).* Platforms (e.g., airplanes) and munitions (e.g., missiles) are increasingly

²⁶ Second Incident of Cyber-Terrorism in Sri Lanka, available at <http://www.lankaweb.com/news/items01/210501-2.html>.

controlled by microelectronics, and such platforms may be sold or made available to other parties (e.g., friendly nations or insurgent groups). But there may be no assurances that these items will not ever be used against U.S. forces. To guard against this possibility, the electronics of such systems could be programmed to self-destruct if a “stay-alive” code were not entered after a fixed period of time, or if the hardware saw a particular bit stream on a communications bus or in memory. The “self-destruct” bit stream could, in principle, be transmitted by U.S. forces confronted with these platforms or munitions.

- *Degradation of adversary smart munitions and platforms (example 2).* Zendia acquires smart weapons using GPS chips made in a factory in a country friendly to the United States. Unbeknownst to Zendia, the GPS chips have circuitry such that if they are given coordinates within the borders of the United States or its allies, they actually translate the coordinates in a random direction to 2 times the damage radius that the United States has calculated for the weapons in use. The weapons test fine for Zendia on all ranges, and work fine when they are used in a skirmish against a neighbor. However, in any engagement with an U.S. ally, the weapons consistently fail to hit targets, and there is no adjustment possible because of the random nature of the translation.

- *Attacking adversary warfighting or warmaking infrastructure (the adversary defense industrial base).* A cyberattack might be used to gain access to a factory producing electric motors for military vehicles. (The factory in question is poorly managed and produces motors only for military use.) With a few commands, the factory is redirected to produce motors using materials that are badly suited for the demands of heavy military use. Such motors work for a short time, but by the time the problem is discovered, many such motors have been shipped and installed in the adversary’s military vehicles.

3.5.3 Cyberattack in Support of Other Operations

Cyberattack can support a variety of other operations as well, though these are not in the category of what are traditionally undertaken by military forces. Illustrative cyberattacks against terrorist groups or international organized crime are described in Chapter 4, on the intelligence community; illustrative cyberattacks to support cyberexploitation on domestic criminals are described in Chapter 5, on domestic law enforcement.

However, an important point to note is that irrespective of whether the intelligence community or domestic law enforcement agencies find it useful and appropriate to conduct cyberattacks against some adversary, it may well be that the U.S. military is the only U.S. government agency with the technical capacity to launch appropriately focused cyberattacks

of significance. Thus, if U.S. military assets and personnel are needed for such purposes, appropriate interagency understandings would have to be reached—and necessary legal authorities obtained—to allow the DOD to execute cyberattacks on behalf of any of these other agencies.

For illustrative purposes only, the examples below describe how cyberattack might be used in support of non-military objectives:

- The leader of an adversary nation controls significant military forces, presides over significant human rights violations in his country, and enriches himself at public expense. A cyberattack could be one approach to threatening the leader's personal financial assets. The existence of such a personal threat might be useful in influencing the leader to stand down his military forces when peacekeeping forces arrive.

- Cyberattack might be an element of a strategic communications effort with the population of a nation. Just as radio has been used as a medium through which the United States has been able to provide information unfiltered by the governments of nations of interest (e.g., Radio Free Europe), the Internet is such a medium today and for the future. However, since nations have been known to seek to block information flows that they regard as unfriendly, U.S. cyberattacks might be used to help residents of these nations circumvent or avoid these various blocking mechanisms.

- Cyberattack might be an element of a strategic communications effort against an adversary. For example, some terrorist groups are known to use the World Wide Web for recruiting purposes and the Internet for communications. Cyberattacks might be used to compromise recruiting websites or servers known to be used by terrorists. Another scenario relates to a kinetic attack on a nation that is accompanied by a cyberattack against that nation's government and media websites. Such an attack might be used to inhibit that nation's ability to tell the world its side of the story,²⁷ or perhaps even to assume control of those websites and provide the world (and its own citizens) with information more favorable to the attacker's position.

It must be emphasized that the scenarios described above are not endorsed by the committee as being desirable applications—only that

²⁷ According to press reports, a cyberattack on Georgian government websites was launched (perhaps by the Russian government, perhaps by private parties sympathetic to the Russian attack) to coincide with the August 2008 Russian attack on South Ossetia, which had the effect of limiting the Georgian government's ability to spread its message online and to connect with sympathizers around the world during the fighting with Russia. See John Markoff, "Before the Gunfire, Cyberattacks," *New York Times*, August 13, 2008, available at http://www.nytimes.com/2008/08/13/technology/13cyber.html?_r=1&oref=slogin.

they represent kinds of scenarios that arise naturally in discussions about cyberattack in pursuit of large-scale strategic interests. As an illustration of a potential problem with such scenarios, consider that manipulation of the information on the websites of an adversary nation's government might affect the information received by U.S. citizens (e.g., through news media receiving altered or manipulated information from those sources and broadcasting that information in the United States). To the extent that the altered or manipulated information was untrue, the U.S. government might be explicitly responsible for misleading the public—an action that could negatively affect the free speech rights of U.S. citizens.

3.6 OPERATIONAL PLANNING

Operational planning processes for cyberattack are not known publicly. But given the similarities of Air Force doctrine for air operations and the cyber missions laid out in Section 3.1, it is not unreasonable to suggest one notional planning process for cyberattack that is roughly parallel to the process for planning offensive air operations—specifically the development of the air tasking order (ATO) that specifies at a high level of detail the actions of air assets in a specific conflict for a specific period of time (usually, 24 hours). The development of a notional cyberattack tasking order (CTO) might entail the following steps.

- The starting point is the explication of a commander's objectives and guidance, and his vision of what constitutes military success. The intent of the operation is defined, and priorities are set. The commander's intent drives the development of targeting priorities and the appropriate rules of engagement. For example, the commander would determine if the intent of the cyberattack is to create widespread chaos or very specific targeted damage.
- The next step is target development. Subject to requirements imposed by the law of armed conflict and the rules of engagement, targets are nominated to support the targeting objectives and priorities provided by the commander. Targets are selected from a variety of sources, including requests from the field, reconnaissance, and intelligence recommendations. Target development often begins before hostilities begin, and the end product of target development is a prioritized list of targets. Legal issues enter here regarding whether a proposed target is indeed a valid and legitimate military target (the necessity requirement discussed in Chapter 7).
- Then comes weaponing assessment. In these phases, the target list is matched to the appropriate types of weapons in the inventory, taking into account the expected results of using weapons on these tar-

gets. Knowledge of munition effectiveness is thus an essential aspect of weaponeering. Legal issues enter here regarding whether the military value of destroying the target outweighs the collateral damage that might occur during the attack (the proportionality requirement, discussed in Chapter 7).

- Force execution refers to the actual execution of the various forces allocated to servicing the targets on the target list, and is the phase in which all elements of the operation are integrated to gain maximum effect. A cyberattack tasking order could support other combat operations and other combat operations could support cyber operations which could be their principal role. Deconfliction (i.e., coordination of forces to ensure that they do not interfere with each other) is part of force execution. For a cyberattack, two phases of execution may be required. An initial phase may introduce a vulnerability that can be exploited later, though if an exploitable vulnerability already exists, this phase may not be necessary. A later phase (perhaps much later) involves the actual exploitation of the vulnerability to cause the damage desired.
- Combat assessment evaluates the effectiveness of combat operations against the commander's objectives. Combat assessment includes battle damage assessment and recommendations for reattack, and it provides the inputs for the next iteration of the cyberattack tasking order.

Another notional process for operational planning of cyberattack might be similar to that used to develop the Single Integrated Operating Plan (SIOP) for using nuclear weapons.²⁸ It is publicly known that the SIOP contains a variety of options from which the President may select should he decide that nuclear weapons should be used. These options fall into categories such as "Major Attack Options," "Selected Attack Options," "Limited Attack Options," "Demonstration Use," and so on. Any given option consists of a list of targets, a timetable on which the targets are to be attacked, and the nuclear weapons systems that are to be used in the attack on those targets.

Translated into the cyberattack domain, a cyber-SIOP could similarly include a list of targets, a timetable on which the targets are to be attacked, and the cyberweapons that are to be used in the attack on those targets. Large-scale attack options might involve large attacks intended to create far-reaching effects, while small-scale options might be narrowly tailored to address a particular target set. Depending on the rules of engagement and the authorizations needed to execute such a plan, either STRATCOM

²⁸ The name of the strategic nuclear response plan was changed to OPLAN 8044 in early 2003. The SIOP terminology is retained here because it is less cumbersome than OPLAN 8044.

or the geographic combatant command could carry out any one of these options, though it is likely that STRATCOM is largely responsible for planning regional attack options as well as attack options relevant to the entire globe.

A major difference between a cyber-SIOP and a nuclear response plan is the possibility of rapid changes in defensive postures for cyber targets. Many of the targets in any nuclear response plan would be fixed in location, with no defensive measures in place. To the extent that cyber targets might change their defensive postures in ways unknown to a cyberattacker, they are more analogous to targeting mobile assets in the nuclear response plan—and targeting of mobile assets is known to be an extraordinarily challenging task. The operational implication of a cyber-SIOP is that a static planning process is unlikely to be effective, and both intelligence gathering and attack planning on possible targets in the various attack options would have to be done on a frequent if not continuous basis.

3.7 HUMAN CAPITAL AND RESOURCES

As the U.S. armed forces become more involved with offensive cyber operations, it becomes more important to have a professional military corps that is actively engaged in thinking about how best to use the new capabilities associated with cyberattack.

From an operational perspective, the complexity and scope of cyberattack suggest that the mix of skills needed to operate successfully is quite broad. Moreover, the necessary skills are not limited to the traditional military specializations of operations, intelligence, and communications—necessary specialized knowledge and information may be needed from the private sector or from other government agencies (e.g., the State Department or Department of Commerce or the Office of the U.S. Trade Representative).

Thus, the operational planning process must include some ways of making such expertise available to military planners and decision makers. Note also that a distributed planning process is also more logistically cumbersome than one in which all the individuals with relevant expertise are available in one location (and are in the same time zone).

Another problem regarding the specialized expertise brought to bear in operational planning is the highly classified nature of cyberattack. With such classification practices in widespread use, it becomes difficult to gain broad exposure to the techniques and the operational implications of employing those techniques—and thus the available expertise is more restricted than it would otherwise be.

Yet another issue is that, as noted in Chapter 2, the success of a cyber-

attack may well depend on the availability of skilled operators who can think “on the fly” and adapt an attack in progress to circumvent unexpected defenses and unanticipated problems. This fact has many implications for training and suggests the importance of focusing on developing cyberattack skills to a very high level of proficiency in a few individuals in addition to developing basic skills in a large number of individuals.

Today, cyberattack operators do not have their own specialization, and they are often typically drawn from those in the intelligence and communications career tracks. (In other cases, they are drawn from combat specializations that do not nurture any particular expertise relevant to cyberattack at all.) In the long run, the increasing skill requirements described above for conducting successful cyberattacks suggest a need for specialization comparable to the more traditional combat specializations for personnel. Such specialization—likely in operations rather than intelligence or communications—would provide training and education that integrates the relevant skills from all of the relevant disciplines. It would also provide upward mobility and well-defined career paths with opportunities for multiple promotions and senior leadership.

Lastly, the Department of Defense invests heavily in realistic training and exercises for personnel with traditional military specializations. Training and exercises go far beyond developing individual competence and expertise in combat—they are proving grounds for new tactical concepts and provide insight into how groups of people (i.e., units) can function effectively as a team. Today, traditional military exercises may include a cyber component, but often the cyber component is not prominent in the exercise and only a relatively small fraction of the exercise involves cyber activities.

The investment in training and exercises for cyberattack and cyberconflict is far below that which is allocated to training for combat in traditional domains. However, not enough is known to determine if the current investment is adequate (that is, if it properly reflects the importance and scale of cyber operations in the future) or inadequate (as might be the case if institutional pressures and prejudices gave short shrift to this type of combat). As this report was going to press, Secretary of Defense Robert Gates announced that in order to improve cyberspace capabilities, the DOD will seek to increase the number of cyber experts that the department can train from 80 students per year to 250 per year by FY 2011.²⁹

²⁹ “Gates Unveils Overhaul of Weapons Priorities,” *Wall Street Journal*, April 6, 2009, available at http://online.wsj.com/article/SB123904207376593845.html?mod=googlenews_wsj.

3.8 WEAPONS SYSTEMS ACQUISITION

The acquisition of weapons is one of the prime responsibilities of the military services. To illustrate some service desires for cyberweaponry:

- The Air Force is seeking to acquire a Cyber Control System (CCS) to provide command and control for the Air Force portion of the DOD Global Information Grid (GIG). The CCS is intended to enable active defense operations “by providing GIG situational awareness along with both automated responses (based on pre-defined Rules of Engagement) and recommended Courses of Action (COA) in response to network intrusions/attacks.” The CCS is also intended to enable network attack operations.³⁰

- The Air Force is supporting the Dominant Cyber Offensive Engagement problem, which is intended to develop capabilities for gaining access to any remotely located open or closed computer information systems; obtaining full control of a network for the purposes of information gathering and effects-based operation; and maintaining an active stealthy but persistent presence within the adversaries’ information infrastructure.³¹

- The U.S. Air Force has noted a need for new technologies to support network attack (network-based capabilities to destroy, disrupt, corrupt, or usurp information resident in or transiting through networks), network defense (network-based capabilities to defend friendly information resident in or transiting through networks against adversary efforts to destroy, disrupt, corrupt, or usurp it), and network warfare support (actions tasked by or under direct control of an operational commander to search for, intercept, identify, and locate or localize sources of access and vulnerability for the purpose of immediate threat recognition, targeting, planning, and conduct of future operations such as network attack).³² Some of these specific needs are described in Box 3.5.

- The Army has issued a broad agency announcement seeking technologies for network disruption using “subtle, less obvious methodology

³⁰ See <http://www.fbo.gov/spg/USAF/AFMC/ESC/R1739/SynopsisP.html>.

³¹ FUNDING OPPORTUNITY NUMBER: BAA 08-04-RIKA, <https://www.fbo.gov/index?s=opportunity&mode=form&id=b34f1f48d3ed2ce781f85d28f700a870&tab=core&cvview=0&cck=1&au=&cck=>

³² Broad Agency Announcement (BAA ESC 07-0001), OL-AA 950 ELSG/KIS, Network Warfare Operations Capabilities (NWOC), Technology Concept Demonstrations, available at http://www.herbb.hanscom.af.mil/tbbs/R1528/Final_NWOC_BAA_Amend_5.doc.

BOX 3.5 Illustrative U.S. Air Force Technology Needs for Cyberattack

A broad agency announcement from the U.S. Air Force calls for proposals to develop the following technologies for network attack, network defense, and network warfare support.¹ Some of the technologies sought include:

- Mapping of networks (both data and voice);
- Access to networks;
- Denial of service on current and future operating systems and network devices;
- Data manipulation;
- Technologies/concepts for developing capabilities for IO modeling and simulation;
 - Situational awareness that gives the operator near real-time effectiveness feedback in a form that is readily observed by the operator;
 - Technologies/concepts for developing capabilities to assess and visualize non-kinetic effects;
 - Technologies/capabilities/concepts for generating and distributing dynamic electronic target folders to include non-kinetic courses of action (COAs);
 - Processing of multi-level security information; and
 - Technologies/concepts for developing capabilities to support rapid implementation of effects-based capabilities.

¹ Broad Agency Announcement (BAA ESC 07-0001), OL-AA 950 ELSG/KIS, Network Warfare Operations Capabilities (NWOC), Technology Concept Demonstrations, available at http://www.herbb.hanscom.af.mil/tbbs/R1528/Final_NWOC_BAA_Amend_5.doc.

that disguises the technique used and protecting the ability whenever possible to permit future use.”³³

Acquisition policy in general terms is addressed in Chapter 6.

³³ Army Offensive Information Operations Technologies Broad Agency Announcement, May 3, 2007, available at [https://abop.monmouth.army.mil/baas.nsf/Solicitation+By+Number/9BE5D8EAE22A6339852572D4004F0DD5/\\$File/BAA+Army+Offensive+Information+Operations+Technologies.doc](https://abop.monmouth.army.mil/baas.nsf/Solicitation+By+Number/9BE5D8EAE22A6339852572D4004F0DD5/$File/BAA+Army+Offensive+Information+Operations+Technologies.doc).

4

An Intelligence Community Perspective on Cyberattack and Cyberexploitation

The intelligence community's primary role relates to the process of generating finished intelligence from raw information for policy makers to use in decision making and action regarding national security and foreign policy. In addition, as a matter of policy and in accordance with legislation and executive order, the Central Intelligence Agency has an operational role in undertaking covert action intended to influence events abroad.

The reader should keep in mind that this chapter is necessarily less complete than the discussion of Chapter 3 since much less is known publicly about the intelligence community's thinking about cyberattack and cyberexploitation. Furthermore, all of the scenarios described below are entirely hypothetical.

4.1 INTELLIGENCE COLLECTION AND ANALYSIS

4.1.1 Governing Principles

In the domain of national security, intelligence is useful for both tactical and strategic purposes. Tactical intelligence is useful to the military services, because it provides advantages on the battlefield against adversary forces through direct support to operational commanders in areas such as reconnaissance, mapping, and early warning of adversary force movements or other actions. Tactical intelligence is also necessary for counterterrorism efforts that seek to preempt or disrupt terrorist activi-

ties before they occur. Intelligence for strategic purposes (national intelligence) serves foreign policy, national security, and national economic objectives. National intelligence focuses on foreign political and economic events and trends; strategic military concerns such as plans, doctrine, and scientific and technical resources; weapons system capabilities; and nuclear program development.

The intelligence-generation process, usually described as a cycle, has several steps. It begins with *planning and direction*, which identifies decision-maker needs for information about a potential adversary (or perhaps even a friendly party). These needs constitute the basis for information *collection* requirements, which specify the scope and nature of the raw information that may be needed in analysis. As a rule, information can be collected from many sources, including open sources such as foreign broadcasts, newspapers, periodicals, books, and websites. Other sources of information are secret, and may include agents abroad, defectors from adversaries, or information clandestinely gleaned from telephone, radio, or Internet transmissions. Information *processing* converts the large amounts of raw information into forms usable by intelligence analysts, and may entail decryption, language translations, and data reduction. *Analysis and production* converts information into finished intelligence and involves integrating, evaluating, and analyzing all available information from all sources. Such analysis may take place over the course of days or weeks or months (in the case of strategic intelligence) or over the course of hours or minutes (in the case of tactical intelligence). *Dissemination* distributes the finished intelligence to the decision makers who requested the intelligence in the first place. (The cyclical nature of the intelligence process results from the fact that recipients of intelligence often develop new requirements and intelligence needs after they receive finished intelligence, and the cycle starts anew.)

The information collection step is the most relevant to this report. Traditionally, sources of information have included signals intelligence (SIGINT—information derived from intercepted communications, radar, telemetry, and computer networks), imagery (IMINT—overhead and ground imagery), measurement and signature intelligence (MASINT—technically derived intelligence data other than imagery and SIGINT, examples of which might be the distinctive radar signatures of specific types of aircraft or the composition of air and water samples), human-source intelligence (HUMINT—including clandestine source acquisition of information; overt information collection by civilian and military personnel assigned to U.S. diplomatic and consular posts; debriefing of foreign nationals and U.S. citizens who have traveled abroad or have access to foreign information; official contacts with foreign governments, including liaison with their intelligence and security services), and open-source

information (OSINT—publicly available information appearing in print or electronic form).

In the context of this report, activities generally labeled as exploitation are sources of raw information and support the information collection step of the intelligence cycle. As noted in Chapter 1, exploitation operations use adversary information systems and networks to support friendly goals and missions.

Computer-based or network-based exploitation operations can be used to support information collection, although they do not necessarily fit neatly into any one of the several sources described above. For example, software agents can be introduced into a collection target's computer system that can scan all accessible files for certain keywords (e.g., "nuclear" in the appropriate local language) and e-mail those files in encrypted form to an address controlled by U.S. intelligence services. Other types of agents can monitor all keystrokes made on a target's computer keyboard. A hardware agent introduced during the design of a microprocessor might secretly render its encryption functions useless for practical purposes, thus making eavesdropping on encrypted messages from that computer relatively easy to perform.¹

Finally and as noted in Chapters 2 and 3, cyberattack often requires substantial intelligence support to succeed, and often cyberexploitation techniques will be used to acquire such information for this purpose. Intelligence agencies of the U.S. government will play a significant role in collecting the intelligence information necessary for such operations by the U.S. armed forces.

4.1.2 How Cyberexploitation Might Be Used to Support Intelligence Collection

Some tools for intelligence collection are based on the clandestine installation of a software or hardware agent into an adversary computer system or network. Once installed, the functionality of the agent for intelligence collection depends only on its ability to route information back to its controller, however circuitous or opaque that route might be.

The following hypothetical scenarios may be illustrative:

¹ Famed cryptographer Adi Shamir noted that "if some intelligence organization discovers (*or secretly plants*) [emphasis added] even one pair of integers a and b whose product is computed incorrectly (even in a single low-order bit) by a popular microprocessor, then ANY key in ANY RSA-based security program running on ANY one of the millions of PCs that contain this microprocessor can be trivially broken with a single chosen message." See Adi Shamir, "Research Announcement: Microprocessor Bugs Can Be Security Disasters," November 2007, available at <http://cryptome.info/bug-attack.htm>.

- The director of the Zendian intelligence service is known to be a strong supporter of the Zendian national soccer team. The soccer team maintains a website on which it provides team statistics, video highlights from recent games, and other content of interest to fans. An intelligence collection operation is launched to exploit a flaw in the operating system of the server that handles the soccer team's website, and installs a Trojan horse program as a modification of an existing videoclip. When the director views the clip, the clip is downloaded to his hard drive, and when his desktop search program indexes the file, the Trojan horse is launched.² The collection payload then searches the local hard drive for evidence suggesting that the user is in fact the director. If none is found, the program erases itself. If the program finds evidence that the user is the director of intelligence (or perhaps the minister of defense, also known to be a soccer fan), it retrieves all plaintext files within reach and e-mails encrypted compressed versions of them to an e-mail address set up specifically as a "dead-drop" location.

- The Zendian Secret Internet Protocol Router Network (Z-SIPRNet) carries classified information and messages for the Zendian ministry of defense, and supports the Zendian command and control system for managing troop deployments, the Zendian defense message system, and many other classified warfighting and planning applications. Although no connections between Z-SIPRNet and the public Internet are allowed, it is known that Gorga, a system administrator, has connected his computer at work to a password-protected dial-up modem. Through a manipulation of the telephone switching center, phone calls from Gorga's home phone number to the modem are secretly redirected to a login simulator that captures his login name and password. Using Gorga's administrator privileges, the intelligence collection operation installs a "sniffer" on the network that examines all passing traffic, and forwards interesting communications to a file that is saved in a temporary work area on Gorga's computer. At night, while Gorga is asleep, the collection operation downloads the file.

- An intelligence collection operation scatters inexpensive universal serial bus (USB) flash drives in parking lots, smoking areas, and other areas of high traffic near a building associated with the Zendian

²For example, a vulnerability in the way in which Windows operating systems handled Windows Metafile vector images was reported in late 2005—this vulnerability allowed arbitrary code to be executed on any computer affected without the knowledge or permission of its users upon viewing of certain image files. See Swa Frantzen, *WMF FAQ*, January 7, 2006, available at <http://isc.sans.org/diary.html?storyid=994>.

Ministry of Defense.³ In addition to some innocuous images, each drive has already-loaded software that collects passwords, login names, and machine-specific information from the user's computer, and then e-mails the findings to the intelligence collectors. Because many systems support an "auto-run" feature for insertable media (i.e., when the medium is inserted, the system automatically runs a program named "autorun.exe" on the medium) and the feature is often turned on, the intelligence collectors can receive their findings as notified as soon as the drive is inserted. The program also deletes itself and any trace of the e-mail after sending. The login information can then be used to compromise the security of existing accounts.

- A Zendian firm and a Ruritanian firm are competitors for a multibillion-dollar contract in a third country. Working closely with the Zendian firm to understand what it would need to know to compete more effectively, the Zendian intelligence service conducts against the Ruritanian firm a series of cyber offensive actions that install dual-purpose and well-hidden Trojan horses on the firm's network. At first, these Trojan horses are programmed to send back to Zendian intelligence confidential business information about the Ruritanian bid; this information is subsequently shared with the Zendian negotiating team. Later, as the deadline for each side's best and final bid approaches, the second function of the Trojan horses is activated, and they proceed to subtly alter key data files associated with the Ruritanian proposal that will disadvantage the firm when the proposals are compared side by side.⁴ (Note that these cyber offensive actions combine cyberexploitation with the installation of a capability for subsequent cyberattack.)

In each of these cases, the installed agent copies files (or parts thereof) and then transmits them to the handler. But any access to copy a file could almost as easily rewrite the file with different data, and on many systems do so without evidence. Such an action would convert the intelligence collection agent into a destructive agent as well.

It should be noted that some of the activities in these scenarios would raise legal and policy questions for U.S. intelligence agencies if they were to engage in such activities. These agencies surely possess the technical capability to engage in such activities, but by policy, the United States does not target intelligence assets for the specific purpose of enhancing

³ This exploit is based on an actual experiment reported in 2006. In this experiment, over 75 percent of the drives distributed resulted in a system penetration. See Steve Stasiukonis, "Social Engineering, the USB Way," *Dark Reading*, June 7, 2006, available at http://www.darkreading.com/document.asp?doc_id=95556&WT.svl=column1_1.

⁴ The use of national intelligence agencies to aid private companies is not unprecedented, as noted in Section 2.6.2.

the competitive position of U.S. industries or specific U.S. companies. If it did, U.S. companies might be able to obtain competitively useful and proprietary information about the future generations of foreign products, such as airplanes or automobiles, or about business operations and contract negotiating positions of their competitors.

A potential legal question arises in the action of the U.S. government in conducting a cyber offensive action against any viewer of a given website, which could include U.S. citizens. Section 7.3.4 addresses the legality of such actions taken by intelligence agencies against foreign or domestic computers, but additional uncertainties arise if such activities are regarded as infringing on the constitutional rights of U.S. citizens.

4.2 COVERT ACTION

4.2.1 Governing Principles

By law (50 USC 413b(e)), covert action relates to activities of the U.S. government to influence political, economic, or military conditions abroad, where it is intended that the role of the U.S. government will not be apparent or acknowledged publicly. Covert action must support identifiable foreign policy objectives of the United States and be important to the national security of the United States, and must be authorized by findings of the President. Covert action must not violate the Constitution or any statute of the United States, nor influence United States political processes, public opinion, policies, or media, and must also be appropriately reported to appropriate individuals in the U.S. Congress. (The legal basis for covert action is addressed in greater detail in Chapter 7.)

In general, covert action is not focused primarily on activities related to intelligence collection or analysis, although such collection may occur incidentally to covert action. Executive Order 12333 stipulates that the Central Intelligence Agency has by default the lead role in covert action.

Classic examples of covert action include providing weapons or funding to a favored party in a conflict, supporting agents to influence political affairs in another nation, engaging in psychological warfare, disseminating disinformation about a disfavored party, or deceiving a disfavored party. Specific actions that could be undertaken under the rubric of covert action include:

- Funding opposition journalists or newspapers that present negative images of a disfavored party in power;
- Paying intelligence agents or party members to make public statements favorable to U.S. interests;

- Providing financial support to opposition civil society groups and helping them set up international networks;
 - Advancing conditions for economic disruption by creating fuel shortages, promoting hoarding, making doomsday predictions, or closing key markets;
 - Providing military aid or training to favored parties;
 - Bolstering individual leaders favorable to the United States who could plausibly fill a power vacuum once the party in power is ousted;
 - Funneling money to a favored party through legal or illegal means;
 - Supporting paramilitary action against a disfavored government of a foreign nation;
 - Instigating a fight or discord between two adversarial, disfavored parties;
 - Influencing an election; and
 - Disseminating propaganda.

As a practical matter, the findings process of the covert action statute was established to provide safeguards in situations where the United States would be drawn further into some conflict or the lives of people on the ground were at risk. The “feel and character” of such situations are significantly different from actions such as remotely placing Trojan horse programs in the operating system of a foreign defense ministry—and it would be less likely for decision makers to believe that findings would be necessary to authorize such actions. Nevertheless, covert action—whether it involves computers or not—is subject to the findings and notification process specified by law.

In addition, it is entirely conceivable that activities originally intended to be outside the statutory definition of covert action will evolve over time into such action, at which time the findings mechanism is supposed to be invoked. Put differently, there is a certain threshold (an ill-defined threshold to be sure) that must be met in order to trigger the findings process, and to the extent that an activity remains below or outside that threshold, the safeguards described in the previous paragraph are not operative.

According to Jeff Smith, former general counsel to the Central Intelligence Agency (1995-1996), traditional U.S. interpretations of the laws of armed conflict (LOAC; further described in Chapter 7) require covert action, whether or not it involves violent activities, to be conducted consistent with LOAC’s requirements. (For example, the War Crimes Act (18 U.S.C. 2441) is applicable to all U.S. nationals.) Smith further noted that observance of the spirit and letter of LOAC is generally helpful in any operation in which it is desirable to win the hearts and minds of the people of the nation involved, and in any case increases the likelihood

that other nations will support (or at least less strenuously oppose) U.S. actions.

This discussion of covert action should not be construed as supporting or opposing the notion of covert action in the first place, and a number of points must be kept in mind. First, covert action is predicated on the assumption that the policy goals being supported are indeed sound and appropriate. No covert action can turn bad policy into good policy, even when decision makers are tempted to use covert action to rescue failed policy. In the latter case, it is easy for covert action to *become* the policy, and for decision makers to forget or downplay the original policy goals. Second, covert action is undertaken on the assumption that its link to the U.S. government can be kept secret. Although experience demonstrates that covert action can indeed be kept secret under some circumstances, decision makers cannot assume that any given covert action will be kept secret—and this holds as well for any covert action that might be based on cyberattack capabilities.

4.2.2 How Cyberattack Might Be Used in Covert Action

One alleged U.S. activity involving cyberattack in a covert action occurred in 1982.⁵ According to Thomas Reed, a former National Security Council official, the United States doctored software that was subsequently obtained by the Soviet Union in its efforts to obtain U.S. technology.⁶ At the time, the United States was seeking to block Western Europe from importing Soviet natural gas. The intent of U.S. doctoring was “to disrupt the Soviet gas supply, its hard currency earnings from the West, and the internal Russian economy,” and to support this goal, “the pipeline software that was to run the pumps, turbines, and valves was programmed to go haywire after a decent interval to reset pump speeds and valve settings to produce pressures far beyond those acceptable to pipeline joints and welds.” Soviet use of the doctored software allegedly caused a large explosion in a Siberian natural gas pipeline.

The following additional (and entirely hypothetical) examples of how cyberattack might be used in covert action are presented for discussion only and without comment on the merits of the underlying goals:

⁵ However, since the U.S. statute defining covert action was not signed into law until 1991, it is unclear whether the 1982 action should be considered a covert action in the legal sense of the term.

⁶ Thomas C. Reed, *At the Abyss: An Insider's History of the Cold War*, Ballantine Books, New York, 2004.

- An election is to be held in Zendia, and the predicted margin of victory between the favored and disfavored parties is relatively small. This election will be the first Zendian election to use electronic voting, and the Zendian election authorities have obtained electronic voting machines to administer this election from Ruritania. U.S. intelligence operatives intercept the CD-ROM containing a software update from the Ruritanian vendor en route to Zendia, and substitute a new CD-ROM in the package containing the original update plus additional functionality that will tilt the election toward the favored party.

- A disfavored party is in power in Zendia, and the U.S. government wishes to weaken it. U.S. intelligence operatives conduct a cyberattack against the Zendian Social Services Agency by compromising employees of the agency, using the USB flash drive technique described above. Obtaining access to the Social Services Agency databases, the United States corrupts the pension records of many millions of people in the country. In the next election, the disfavored ruling party is voted out of office because of the scandal that resulted.⁷

- Two traditionally adversarial nations are armed with nuclear weapons, and the United States has been conducting intelligence collection operations against these nations for many years. Through a mix of human and technical means, it has been successful in learning about cyber vulnerabilities in the nuclear command and control networks of each nation. During a crisis between the two nations in which both sides have launched conventional kinetic attacks against the other side's territory and armed forces, nuclear confrontation between them is imminent. The U.S. government makes a decision to corrupt the transmission of any nuclear launch orders transmitted through those networks in order to prevent their use.⁸

- Zendia is an authoritarian nation that recognizes the value of the Internet to its economy, but as an instrument of political control, it actively censors certain kinds of Internet content (e.g., negative stories about the Zendian government in the foreign press) for its population. Its censor-

⁷ This scenario is based on the Japanese election in 2007, in which the ruling party lost resoundingly. Many analysts attributed the loss to the fact that the Japanese Social Insurance Agency was revealed to have lost pension records for 50 million people. Although no evidence suggests that cyberattacks played any role in this scandal, it is easy to see how in an age of increasingly automated records, such attacks might well have such a large-scale effect. See Pino Cazzaniga, "Election Defeat Marks Abe's Political Future," *AsiaNews.it*, July 30, 2007, available at <http://www.asianews.it/index.php?l=en&art=9962>.

⁸ In 1996, a scenario with many similar elements involving India and Pakistan was proposed by John Sheehan, then-commander-in-chief of the U.S. Atlantic Command. See Bradley Graham, "Cyberwar: A New Weapon Awaits a Set of Rules," *Washington Post*, July 8, 1998, p. A1.

ship mechanisms are largely automated and operate at one of a few Internet gateways to the country. During a time of tension with Zendia, the United States launches a cyberattack against the automated Zendian censors so that the population can obtain, at least temporarily, a broader range of information than it would otherwise be able to access.

- A party favored by the United States is conducting an armed rebellion against the Zendian government. No funds are currently available to help the favored party. However, the U.S. President wishes to find a way to help the rebels, and authorizes a cyberattack that diverts money from the Zendian national treasury to the rebels.

- A Zendian cyberattack is launched against the military medical services of Ruritania to obtain the medical records of all active personnel. In the days before a planned armed attack by Zendia, postings and mailings from anonymous sources appear pointing out that Ruritanian Colonel X is being treated for bipolar disorder, that Captain Y was treated three times for a sexually transmitted disease in the last 2 years, and that Admiral Z is on tranquilizers. Copies of the medical records—sometimes secretly and undetectably altered—were released to back up the stories. The results led to some family problems for Captain Y, Admiral Z was relieved of field command, and Colonel X resigned his commission. Others were simply discomfited. The result was a drop in readiness by the command structure when Zendia struck, giving Zendia some advantage. Note that this particular covert action has an element of intelligence collection.

- The Zendian nuclear weapons program relies on a social network of scientists and engineers. The United States launches cyberattacks against a dozen key scientific leaders in this network to harass and discredit them. These cyberattacks plant false adverse information into their security dossiers, insert driving-under-the-influence-of-drugs/alcohol incidents into their driving records, alter their credit records to show questionable financial statuses, change records of bill payments to show accounts in arrears, and falsify telephone records to show patterns of contact with known Zendian criminals and subversives.⁹ Discrediting these individuals throws the program into chaos.

- Scientists working on the Zendian biological weapons program use an in-house network to communicate with each other and manage their research and development program. U.S. intelligence agencies penetrate the network to install dual-purpose software agents to exfiltrate the traffic on the network to intelligence analysts. When analysis of the traffic indicates that the Zendian research efforts are reaching a critical stage,

⁹ This scenario is based on one taken from the Global Organized Crime Project, *Cybercrime, Cyberterrorism, Cyberwarfare: Averting an Electronic Waterloo*, Center for Strategic and International Studies, Washington, D.C., 1998.

the software agents begin to alter key data clandestinely so that critical experiments fail. Further, these software agents are so well hidden that they can maintain their presence over a period of years so that subsequent experiments fail at critical times as well.

- The Zendian airplane industry and a major U.S. defense contractor are engaged in a competition to win a lucrative contract from Ruritania for producing fighter aircraft. In order to support a key company in the U.S. defense industrial base, the U.S. government conducts a cyberattack to disrupt and delay the production of the Zendian fighter plane and thereby provides an additional incentive for Ruritania to select the U.S.-produced plane.¹⁰

4.3 POSSIBLE INTELLIGENCE COMMUNITY INTEREST IN CYBERATTACK AND CYBEREXPLOITATION

Because such information would fall into the category of sensitive “sources and methods,” it is not publicly known whether the intelligence community has used or intends to use cyberexploitation. However, the use of cyberexploitation techniques for exfiltration of sensitive business and personal information is well known, and the U.S. government has indicated that DOD systems have been subjected to foreign cyberexploitation for such purposes. Thus, it would be highly surprising if the U.S. intelligence community did not know about and make use of cyberexploitation when appropriate or helpful.

As for covert action, again the CIA’s interest in or use of cyberattack is not known publicly. But given the demonstrated difficulties in tracing the source of a destructive cyberattack to a specific party, it would not be at all surprising for the CIA to be interested in cyberattack as at least a potential tool for covert action.

Hints of possible interest in the value of cyberattack for the intelligence community can be found in the testimony of Director of National Intelligence J. Michael McConnell to the Senate Select Committee on

¹⁰ Although such actively destructive actions have not, to the committee’s knowledge, been taken to benefit U.S. companies, U.S. intelligence has been used to uncover unfair trade practices of other nations whose industries compete with U.S. businesses, and has helped the U.S. government to ensure the preservation of a level economic playing field. According to the National Security Agency, the economic benefits of SIGINT contributions to U.S. industry taken as a whole have totaled tens of billions of dollars over the several-year period prior to 1996. See National Research Council, *Cryptography’s Role in Securing the Information Society*, National Academy Press, Washington, D.C., 1996, Chapter 3.

Intelligence in February 2008.¹¹ McConnell noted a need for the United States “to take proactive measures to detect and prevent intrusions from whatever source, as they happen, and before they can do significant damage.” He also noted concern about “how best to optimize, coordinate and deconflict cyber activities.” The first statement points to the inadequacy of hardening and passive defense alone as defensive strategies, and the second statement about coordination and deconfliction suggests the existence of (or the desire to conduct) cyber activities outside one’s own defensive perimeter that might contribute to defense.

Finally, as noted in Box 3.2, the National Security Agency—which is a member of the intelligence community and also a component of the Department of Defense—has in its latter role certain responsibilities for cyberattack activities.

¹¹ J. Michael McConnell, “*Annual Threat Assessment of the Director of National Intelligence for the Senate Select Committee on Intelligence*,” February 5, 2008, available at <http://intelligence.senate.gov/080205/mcconnell.pdf>.

Perspectives on Cyberattack Outside National Security

As noted in Chapters 3 and 4, the military and intelligence communities have missions to which cyberattack capabilities are relevant. But cyberattack may be relevant to at least two other constituencies—the domestic law enforcement community and the private sector. This chapter explores some of those possible connections.

5.1 CYBERATTACK AND DOMESTIC LAW ENFORCEMENT

For many years, the law enforcement community has had the authority to undertake covert surveillance and monitoring of electronic computer-based communications under legally authorized circumstances. (The legal authority for such activity is Title III of the Omnibus Crime Control and Safe Streets Act of 1968, as amended to include the Electronic Communications Privacy Act, and briefly described in Chapter 7.) In addition, law enforcement authorities may conduct surreptitious searches of computers for documents when so authorized under a court-issued warrant.

From a technological standpoint, such activities are equivalent to the intelligence collection activities described in Chapter 4. Law enforcement authorities can and do conduct cyberexploitation with the appropriate legal authorization, although the legal framework for providing authorization is very different for the law enforcement community than for the intelligence community.

By contrast, law enforcement authorities often eschew cyberattack.

One possible reason is the prosecutorial focus of law enforcement authorities, who are generally concerned with obtaining legally admissible evidence in order to support successful prosecution. Evidence collected from the computers of suspects has been subject to claims that computer records have been altered.¹ Absent specific evidence that tampering has occurred, such claims have not prevailed to date. But if an operation were specifically *designed* to damage or destroy information resident on a target computer, it is hard to imagine that such claims would not be taken more seriously.

A second reason may be that other tools are often available. For example, a criminal website in the United States being used to defraud consumers, for example, can be taken down by legal rather than technical means.

On the other hand, public reports indicate that law enforcement authorities have in fact conducted denial-of-service attacks against wireless (cell phone) networks and other wireless devices such as garage door openers and remote control devices for toys in order to prevent their use to detonate remote-controlled bombs.² Jamming cell phone networks in a specific geographic area could be used to help stop terrorists and criminals from coordinating their activities during a physical attack and prevent suspects from erasing evidence on wireless devices. In prisons, jamming could interfere with the ability of prison inmates to use contraband cell phones, which are often used to intimidate witnesses, coordinate escapes, and conduct criminal enterprises.

Federal law enforcement officials are permitted to use jamming technology with specific legal authorization, and state and local law enforcement agencies are not allowed to do so at all. In particular, 47 USC 333 states that “no person shall willfully or maliciously interfere with or cause interference to any radio communications of any station licensed or authorized by or under this chapter or operated by the United States Government.” However, Section 305 of the Communications Act of 1934 (today 47 USC 305) stipulated government-owned radio stations need not adhere to rules and regulations designed to prevent interference with other radio stations. The National Telecommunications and Information Adminis-

¹ U.S. Department of Justice, *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*, Computer Crime and Intellectual Property Section, Criminal Division, July 2002, available at http://www.usdoj.gov/criminal/cybercrime/s&smanual2002.htm#_VB1_.

² Specifically, the *Washington Post* reported that such jamming technology was used to protect President Obama’s inaugural motorcade on Pennsylvania Avenue. See Spencer S. Hsu, “Local Police Want Right to Jam Wireless Signals,” *Washington Post*, February 1, 2009, p. A02, available at http://www.washingtonpost.com/wp-dyn/content/article/2009/01/31/AR2009013101548_pf.html.

tration Organization Act of 1992, Public Law 102-538 and codified at 47 USC 901-904, established the NTIA as the federal point of responsibility for managing U.S. domestic use of the spectrum.³ Specifically, the Office of Spectrum Management within the National Telecommunications and Information Administration of the Department of Commerce develops and implements policies and procedures for domestic issues regarding the use of the spectrum by the federal government in the United States.

The use of jamming technology is not the only way to thwart the use of cell phones for terrorist or criminal purposes. Persuading cell phone providers to shut down service, either over a broad area or just in the vicinity of a few specific cell towers, can also work effectively—such an approach to a cell phone provider might well be regarded as the equivalent of a close-access “cyberattack.”⁴

In February 2009, Senator Joseph I. Lieberman planned to introduce legislation that would give law enforcement agencies “the tools they need to selectively jam” communications in the event of a terrorist attack.⁵ Senator Kay Bailey Hutchison and Representative Kevin Brady also introduced a bill that would allow the U.S. Bureau of Prisons and governors to seek the authority to jam cell phones in prisons.⁶

5.2 THREAT NEUTRALIZATION IN THE PRIVATE SECTOR

5.2.1 Possible Response Options for Private Parties Targeted by Cyberattack

In general, a private party that is the target of a cyberattack has four options for responding. First, it can implement passive measures to strengthen its defensive posture. For example, it can drop functionality on its own systems that the attacker is trying to exploit, reject traffic, and close ports on its firewall. Second, it can report the event to law enforcement authorities, and law enforcement authorities can take appropriate action to try to shut down the cyberattack (e.g., by finding the perpetrator

³ See <http://www.ntia.doc.gov/osmhome/roosa8.html>.

⁴ Indeed, the *Washington Post* story reported that the U.S. Department of Homeland Security reached an agreement in 2006 with cell phone companies to voluntarily shut down service under certain circumstances, which could disable signals for areas ranging from a tunnel to an entire metropolitan region.

⁵ See Spencer S. Hsu, “Local Police Want Right to Jam Wireless Signals,” *Washington Post*, February 1, 2009, p. A02, available at http://www.washingtonpost.com/wp-dyn/content/article/2009/01/31/AR2009013101548_pf.html.

⁶ Matthew Harwood, “Bill Would Allow Prisons to Jam Cell Phone Signals,” *Security Management*, January 16, 2009, available at <http://www.securitymanagement.com/news/bill-would-allow-prisons-jam-cell-phone-signals-005082>.

and arresting him). (However, to the best of the committee's knowledge, law enforcement authorities have never launched a counter-cyberattack.) Third, it can take self-help measures to further investigate and characterize the source of the cyberattack and then report the information to appropriate law enforcement authorities. Fourth, it can take actions to neutralize the incoming cyberattack.

The first option—strengthening its defense posture passively—entails a minimum of controversy as a matter of law and policy. But although stronger passive defensive measures are unlikely to be effective over the long run, the other options do entail some degree of controversy.

Consider the long-standing thread of policy that law enforcement authorities have a key role to play in responding to a cyberattack against a private sector entity. The law enforcement paradigm is oriented primarily toward investigation, prosecution, and conviction of those who violate existing criminal laws about causing damage or destruction to computer systems (described in more detail below). Such processes take time to operate, often weeks or months, and are often constrained by the availability of law enforcement resources and expertise.

In the meantime, the private sector entity subject to a hostile cyberattack can only hope that passive defense measures will mitigate the threat—today, there are no legal mechanisms or institutional structures available to provide immediate relief under such circumstances. Such a lacuna raises the possibility that some form of active defense for threat neutralization (active threat neutralization for short) may be a necessary component of a strong cybersecurity posture for the private sector.

As noted in Chapter 3, the U.S. Strategic Command (STRATCOM) asserts the authority to conduct response actions, including threat neutralization, on behalf of cyberattacks against DOD installations under certain circumstances. The Department of Homeland Security has the responsibility for seeing to the cyber protection of the defense industrial base and the providers of critical infrastructure. But to the best of the committee's knowledge, neither DHS nor any other part of government has been given the authority to conduct active threat neutralization on behalf of any part of the private sector (including the companies of the defense industrial base and the providers of critical infrastructure).

This state of affairs is problematic for large multinational corporations that face major cybersecurity threats, and that are themselves concerned with how to manage the risk associated with the cyberattacks they face. For such entities, one element of any rational risk management strategy would involve managing the tradeoff between the legal liabilities associated with actions for the defense of property and the benefits from taking such actions.

5.2.2 Self-defense by Private Parties

If passive defensive measures by themselves are insufficient for an adequate cybersecurity posture (as might be inferred from the consideration of active threat neutralization for DOD cybersecurity), the question arises as to whether critical parts of the private sector might be afforded a similar kind of protection. Some elements of the private sector with services to offer do make just such an argument.⁷ Without prejudging the pros or cons of such arrangements, the discussion below indicates some of the legal and policy issues that would need to be addressed before such practices could be adopted.

A first point is whether cyberattack expertise is available to the private sector. Although it is likely that the capabilities of the DOD far exceed those available to the private sector, many private sector companies use penetration testing and “red-teaming” as a way of testing their own cybersecurity postures. Such testing involves hiring a firm or an individual to penetrate one’s own information systems—typically these firms and individuals advertise their services under the label of “ethical hackers” or something similar. The expertise needed to provide these services is roughly the same as that needed to conduct cyberattacks against any other target, and so it is clear that some level of cyberattack expertise is available.⁸ In addition, many private enterprises make use (in some cases, extensive use) of threat intelligence and surveillance capability provided by private companies.

As for the legal dimension, U.S. common law admits certain rights of self-defense and of defense of property in preventing the commission of a crime against an individual or a corporation. (In legal usage, self-defense refers to the defense of one’s self (or others)—a defense of a person. Defense of property is more limited, in the sense that the range of allowable actions for the defense of property is more limited than for

⁷ For example, in a paper entitled “Offensive Operations in Cyberspace” (dated June 1, 2007), the White Wolf Security corporation argued that corporate victims of cyberattack have limited rights to use offensive cyber operations in order to proactively protect their assets and workforce from attacks originating in the United States and in allied non-U.S. nations and that private military companies constitute an emerging base from which to conduct such operations on behalf of any party entitled to conduct them. This paper is no longer online, but it is in the committee’s possession and available in this project’s public access file.

⁸ An important area in which necessary cyberattack expertise may vary according to the kind of target is the expertise needed to conduct social engineering attacks, which by definition involve exploitation of vulnerabilities that are embedded in the particular culture and operating procedures and practices of the target entity. It is almost certainly harder for a U.S. “ethical hacker” to conduct a social engineering attack in Zendia than in the United States, for reasons that might include a lack of knowledge of the Zendian language or of Zendian cultural norms.

certain kinds of self-defense (in particular, for self-defense against lethal attack). But the range of allowable actions for the defense of property is roughly comparable to the range for *non-lethal* self-defense—the use of non-lethal force can be justified in order to defend one’s self against non-lethal attack and to defend one’s property. For hostile cyberattacks, the relevant concept will almost always be defense of property, as cyberattacks against private parties have not usually had lethal intent. Note that self-defense in this context has an entirely different meaning than self-defense in international law, a topic explored at length in Chapter 7.) While individuals are not permitted to engage in revenge or retaliation for a crime (that is, vigilantism is forbidden by law), they are—under some circumstances—entitled to take otherwise-prohibited actions for the purpose of preventing or averting an imminent crime or one that is in progress. Moreover, these rights attach even if specific statutes may not explicitly acknowledge their existence.⁹ Thus, the widely held view that government has a literal monopoly on legitimate use of physical force is simply not true as a matter of common law.

Today, the primary federal law addressing cyberattacks is the Computer Fraud and Abuse Act (CFAA), codified as Title 18, Section 1030. Loosely speaking, this act criminalizes the intentional damaging of any computer connected to the Internet.¹⁰ (A number of state laws have similar provisions and would apply to individuals and corporations within their jurisdiction.¹¹ The CFAA is discussed further in Section 7.3.4.) Although the CFAA contains an explicit exception for law enforcement agencies that undertake the normally proscribed behavior with respect to cyberattack, there is no explicit exception for private parties.

On the other hand, the CFAA was never intended to apply and does

⁹ The Model Penal Code does include exceptions for self-defense and defense of property (*Model Penal Code*, American Law Institute, Philadelphia, 1962, available at http://www.ali.org/index.cfm?fuseaction=publications.ppage&node_id=92). See also Paul H. Robinson, “Criminal Law Defenses: A Systematic Analysis,” *Columbia Law Review* 82:199-291, 1982, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=662043.

¹⁰ Section 1030 of the Computer Fraud and Abuse Act penalizes individuals who “knowingly cause the transmission of a program, information, code, or command, and as a result of such conduct, intentionally causes damage [in excess of \$5,000] without authorization, to a protected computer.” “Protected computers” are defined to include computers “used in interstate or foreign commerce or communication, including a computer located outside the United States that is used in a manner that affects interstate or foreign commerce or communication of the United States.” In short, virtually any computer connected to the Internet falls under the definition of “protected computer.”

¹¹ For example, Section 815.06 of Title XLVI of the Florida Code (entitled “Offenses Against Computer Users”) criminalizes the willful, knowing, and unauthorized access or destruction of a computer or network (among other things). See http://www.leg.state.fl.us/Statutes/index.cfm?App_mode=Display_Statute&Search_String=&URL=Ch0815/SEC06.HTM&Title=-%3E2007-%3ECh0815-%3ESection%2006#0815.06.

not apply to penetration testers—private parties hired (authorized) by a company to test its own defenses.¹² A number of such firms provide such services so that a company can obtain a realistic assessment of its own security posture, and indeed penetration testing is often recommended as one of the best ways of doing so.¹³

A more significant issue is that in light of common law traditions regarding self-defense and defense of property, it is at least possible that a court might find that certain cyberattack actions undertaken in defense of property might be allowable, although whether such actions can stand as an exculpatory rationale for conducting active threat neutralization has not been tested in the courts to date. Even if not, actions taken in defense of property might be a starting point for legislative change if a policy decision is made that such actions involving cyberattack *should* be allowed in certain circumstances.¹⁴

In the context of active threat neutralization of private, non-government computer systems under attack, an interesting question thus arises. To what extent and under what circumstances is self-help a legitimate option for the target of a cyberattack to stop it? Box 2.4 in Chapter 2 describes a spectrum of possible responses to a cyberattack, some of which plausibly count as active defense for threat neutralization.

Security specialists in a private organization are often warned about undertaking efforts to gather information about the perpetrators of a cyberattack against the organization. For example, they are warned against compromising an already compromised machine to insert tracking and collection software to gather such information. Concerns sometimes arise over the possibility that the private organization and/or the security specialists themselves might be subject to civil or even criminal liability for their actions and that their efforts might contaminate evidence should a prosecution occur.

As for more aggressive actions, actions taken in self-defense or for the defense of property are often justified as the only timely response available in exigent circumstances when law enforcement authorities are unavailable at the moment they are needed to prevent a crime—that is, in seconds rather than in the minutes or hours that it often takes law enforcement officials to arrive at the scene of a crime in progress. In the

¹² The CFAA criminalizes only intentional damage caused without authorization.

¹³ National Research Council, *Cybersecurity Today and Tomorrow: Pay Now or Pay Later*, The National Academies Press, Washington, D.C., 2002.

¹⁴ As an example of a policy that would endorse cyberattack as a response to a threat to commercial interests, consider the controversial proposal of Senator Orrin Hatch to “destroy” computers that have repeatedly been involved in the online trading of music and movie files after first providing warnings to the user to refrain from such behavior. See Associated Press, “Hatch Wants to Fry Traders’ PCs,” June 18, 2003, available at <http://www.wired.com/entertainment/music/news/2003/06/59298>.

case of responding to hostile cyberattacks, the enormous number of sites subject to such cyberattacks suggests that sufficient government resources will indeed be unavailable to protect all of them, and law enforcement authorities are often hard-pressed to respond at all, let alone adequately, to cybercrimes in progress.

In the absence of sufficient law enforcement resources, two options are possible—prioritize so that government resources are used to conduct actions to defend people and property only against the most serious threats, and/or allow the attacked parties to conduct such actions themselves.

In large part, a choice between these two options rests on one's view about whether the conduct of offensive activities should be the exclusive purview of government. Very few individuals would be sympathetic to the notion of privatizing the nuclear deterrent force or even battleships or jet fighters. Yet under some circumstances, private parties can and do act with lethal force in order to neutralize an immediate threat to life, and they can act with non-lethal force to neutralize an immediate threat to property.

It is not known how frequently victims of cyberattack take self-help actions. Likely because of concerns about violations of the CFAA, few of those who actually take such actions will report them openly. Yet some anecdotal evidence and personal experience of committee members suggests that the frequency is not zero, and the committee is aware of instances in which attacked companies have indeed conducted denial-of-service counterattacks against the attacking parties, even though such actions have never been acknowledged openly or done in ways that draw attention to them.

One data point on this issue is provided by the *New York Times*,¹⁵ which reported that a worm released in late 2008 known as Conficker has reignited a debate inside the computer security community over the possibility of eradicating the program before it is used, by launching a cyberattack to compromise the worm's controller and direct it to send messages to users warning them of the infection. One cybersecurity researcher working on a counter to the Conficker worm said of such a possibility, "Yes, we are working on it, as are many others. Yes, it's illegal, but so was Rosa Parks sitting in the front of the bus." Others in the cybersecurity research community continue to oppose such an effort to stop the worm because of a concern that such efforts would create even more problems.

If a domestic policy decision is made to allow attacked private-sector

¹⁵ John Markoff, "Worm Infects Millions of Computers Worldwide," *New York Times*, January 22, 2009, available at <http://www.nytimes.com/2009/01/23/technology/internet/23worm.html>.

parties to conduct actions in defense of property,¹⁶ mechanisms can be put into place that depend on other parties whose job it is to defend the interests (life, property) of a possible victim. That is, an individual or a company may hire armed guards to implement self-defense practices or procedures should they become necessary. Regarding cyberattack, the analogous situation might be a company that provides active threat neutralization services that could be called into action when a customer reports being under attack (Box 5.1).

5.2.3 Regulating Self-defense by Private Parties

Some cybersecurity analysts propose letters of marque and reprisal as a model for regulated private cyberattacks to support threat neutralization.¹⁷ Letters of marque and reprisal were originally used by governments to give private parties the authority to take certain actions generally regarded as appropriate only for a nation's military forces—namely to operate and use armed ships to attack and capture enemy merchant ships in time of war. These letters were crafted with a certain degree of specificity to ensure that the actions of the private party did not exceed the intent of the issuing government, and further were never intended to imply that such letters were needed for immediate self-defense.

Although the Paris Declaration Respecting Maritime Law of 16 April 1856 was issued to abolish such private actions, and many nations ratified this declaration, the United States did not and has never renounced the right to do so. Indeed, Article 1, Section 8 of the United States Constitution includes the issuance of letters of marque and reprisal as one of the enumerated powers of Congress.

In the context of privately conducted cyberattacks, letters or licensing could be used to specify the circumstances under which threat neutralization may be performed for the defense of property, the criteria needed to identify the attacking party with sufficiently high confidence, the evidence needed to make the determination that any given cyberattack posed a threat sufficiently severe as to warrant neutralization, and the nature and extent of cyberattacks conducted to effect threat neutralization.

A key issue is the threshold at which it is appropriate to conduct an

¹⁶ Although the policy decision would be domestic, it might well have implications for international law as well. In particular, it is not clear how an explicit decision to allow attacked private sector parties to conduct actions in self-defense or in defense of property would square with the international Convention on Cybercrime (discussed further in Section 7.2.4).

¹⁷ See, for example, Excalibur R&D, "Letter of Marque and Reprisal for Fighting Terrorists," August 20, 2008, available at <http://excaliburrd.com/cs/blogs/excalibur/Excalibur%20Letter%20of%20Marque%20paper%2015%20August%202008.pdf>.

BOX 5.1 A Security Operations Center

Nearly all large organizations face daily a deluge of security inputs from a variety of different systems, platforms, and applications. Usually, these inputs are generated as the result of point solutions distributed over multiple locations and do not adhere to any standards of syntax (they come in different formats, for example) or semantics (they report on different things). Growth in the number of attacks experienced every day, new technologies and rapid expansion, and new regulations and laws increase the burden on systems administrators.

In response, many organizations seek to centralize the management of their security functions in what are often known as security operations centers (SOCs). SOCs track and integrate the multiple security inputs, ascertain risk, determine the targets of an attack, contain the impact of an attack, and recommend and/or execute responses appropriate to any given attack. In some cases, an organization will establish a SOC for itself. In other cases, SOC services are outsourced to a private company that specializes in providing such services.

A SOC is constrained today to provide only passive defensive services when a threat originates outside the perimeter of the organization it serves. But active defense could be provided, legally, if undertaken within the organizational perimeter¹—and it is a matter of policy and law rather than technology that would prevent a SOC from taking similar action against parties outside the organizational perimeter. Of course, if policy and law were established to allow such action, a SOC's actions would be subject to whatever standards and regulatory requirements were part of that policy and law.

¹ As noted in Footnote 10 of this chapter, the Computer Fraud and Abuse Act criminalizes only attacks committed *without authorization*. If a SOC conducts active defense within an organization's perimeter at the behest of that organization, it is acting with authorization.

SOURCE: Adapted from Computer Associates, "Best Practices for Building a Security Operations Center," August 2006, available at http://www.secguru.com/files/papers/best_practices_snoc_white_paper.pdf.

active threat neutralization, and how that threshold is determined. Who determines the threshold? What level of actual damage, if any, must the victim sustain in order to demonstrate harm? How are such levels to be measured? Who should have the authority to make such a determination? What alternatives to active threat neutralization must have been tried before active defense can be used? How should their success (or lack thereof) be documented?

Although in a cyberattack context, these questions reflect largely unexplored legal territory, a few speculations can be made based on past precedents. To be justified, lethal actions taken in self-defense must usu-

ally be carried out as a last resort,¹⁸ but the same is not true for non-lethal actions taken either in self-defense or in defense of property. That is, in the case of protecting one's self from lethal attack, the ostensible victim must have tried other less violent methods for mitigating the harm that an attack might cause, or at least have good cause to believe that those other methods would not work. The same is not true for taking non-lethal actions to combat a non-lethal threat. In the case of cyberattack, even though such attacks are generally non-lethal, an argument might be made that the former analogy imposes a degree of prudence that is appropriate—if so, the analogy would require that a victim has taken all available passive defense measures (e.g., firewalls or system patches) before active threat neutralization is to be allowed as a permissible self-help action.

A related point is that an action in defense of property might be misdirected and thus harm an innocent third party. From a legal point of view, a party taking an action in defense of property resulting in such harm may have a plausible defense to the violation of criminal law if he has made reasonable efforts to identify the party responsible for the original attack, even if the efforts were erroneous. Civil liability may attach for such action (e.g., the party launching the action in defense of property may be responsible to the innocent victim for damages suffered), although the liability might be less if the innocent party was negligent in allowing his or her computer to be used for malevolent purposes.¹⁹

5.2.4 Negative Ramifications of Self-defense by Private Parties

The discussion above should not be construed as advocating a change from today's legal regime that strongly discourages active threat neutralization by private sector entities. Indeed, allowing self-help actions for private parties also has a variety of broader and negative ramifications for the nation's interests writ large.

¹⁸ In most states, it is legal to use deadly force against an attack that threatens death, serious bodily injury, rape, kidnapping, or in some states robbery or burglary, even if one could have safely avoided the injury by retreating. And in all states, it is legal to use deadly force against such an attack even if one could have safely avoided the problem by turning over property that the attacker is demanding as a condition of not injuring the victim.

¹⁹ Uncertainties abound in this area. For example, it is theoretically possible that cyberattack conducted to neutralize an active threat might itself be characterized as an "ultra-hazardous activity," depending on its scope and nature. If so, the courts could apply strict liability to all the harmful effects it caused. Alternately, if the defense of property is found not to apply to an active threat neutralization, the defender could easily find his responsive acts characterized as wrongful. Still other legal traditions forbid "hot pursuit" of an attacker after he no longer poses a threat, and the line between active threat neutralization and retaliation is not necessarily clear.

For example, active threat neutralization conducted by the private sector may have negative implications for the conduct of international relations. A private party in the United States conducting an action that harms computers in Zendia is likely to be attributed to the U.S. government even if there is no such linkage, and Zendia may well seek to hold the United States government responsible. A denial by the U.S. government may even be seen as evidence of government complicity in a plausibly deniable attack. And if Zendia believes that the U.S. government is responsible for an attack on it, it—or computer-savvy citizens of Zendia—may well see fit to attack the United States directly or its interests (e.g., private sector companies). (Such complex escalation scenarios do not generally characterize the typical self-defense or defense-of-property scenarios of a company defending its building or a homeowner defending her home.)

In addition, active threat neutralization conducted by the private sector may also interfere with cyberattacks launched by the U.S. government. For example, it is easy to imagine a scenario in which major U.S. corporations come under cyberattack from a foreign power and the U.S. government chooses to respond in kind. Cyberattacks launched by these corporations at the same time might well interfere with the conduct of the U.S. cyberattack, might work at cross-purposes with it, and would almost certainly be indistinguishable from cyberattack actions taken by the U.S. government.

These issues are further complicated if the U.S. government establishes standards, mandates licensing, or otherwise provides advice that could support actions taken in defense of property (e.g., that describe what conditions must be established for when such behavior should be considered a reasonable option, or what the limits on such actions should be).

In the absence of mandatory standards for taking such action, actions by private parties would be governed by the party's own view of its self-interest, and in particular would be unlikely to take into account other broader societal or national needs.²⁰ Thus, active threat neutralization may run a higher risk of having effects that work against those broader needs or objectives. A private party's threshold for action may also be lower (for example, it may be less tolerant of corporate espionage) than public policy might dictate.

²⁰ Precedent for this likely outcome can be found in the behavior of private companies today, which invest in cybersecurity to the extent, and only to the extent, that their business needs call for such protection. The result has been a societal level of cyber protection that is, by most accounts, inadequate to meet the needs of the nation as a whole. See National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

On the other hand, the explicit establishment of stated policy that allowed private parties to act in defense of property to a cyberattack could well be taken as government endorsement of such actions, even if such policy did not require them to do so. Standards established ostensibly to regulate such behavior and prevent these actions from being taken arbitrarily or solely at the discretion of the victimized party could thus have a perversely negative effect on how the U.S. government is perceived.

Self-help actions of multinational corporations have implications with respect to both international law and the domestic laws of all the nations in which the corporations have a physical presence (where, for example, personnel and assets may be placed at risk by actions taken elsewhere by the corporation). Although such actions have not yet produced a visible reaction from other nations (perhaps because the scale of the problems involved has not reached the necessary level), how nations and the international community will react in the future remains to be seen.

Some of the negative ramifications described are also associated with today's regime, in which victims sometimes take self-help actions on the basis of their own judgments and perceptions quietly and under the table without policy or legal guidance. If and when such self-help actions reach a level where they interfere significantly with U.S. policy or its execution, policy makers may eventually consider a legal regime that is tighter with respect to self-help rather than looser than that of today. A tighter regime might explicitly prohibit active threat neutralization by private parties even under the rubric of defense of property, prohibit active intelligence gathering by private parties in the wake of a cyberattack, make parties undertaking threat neutralization strictly liable for any harm they cause, and so on.

5.3 CYBEREXPLOITATION IN THE PRIVATE SECTOR

Given that the technical skills for cyberexploitation are similar to those required for cyberattack and in light of the discussion above, it is likely that some U.S. companies would have the technical capability to conduct cyberexploitation against their competitors. However, the Economic Espionage Act of 1996, 18 USC 1831-1839, criminalizes the theft of trade secrets related to or included in a product that is produced for or placed in interstate or international commerce. ("Trade secrets" are defined as "all forms and types of financial, business, scientific, technical, economic, or engineering information, including patterns, plans, compilations, program devices, formulas, designs, prototypes, methods, techniques, processes, procedures, programs, or codes.") Whether individual U.S. firms have engaged in this kind of industrial intelligence against competitors despite its illegality is unknown to the committee.

5.4 THREAT NEUTRALIZATION ON BEHALF OF NON-MILITARY GOVERNMENT AGENCIES

Most of the discussion in the previous sections applies equally as well to active threat neutralization conducted on behalf of non-military government agencies. The Department of Homeland Security has the responsibility for seeing to the cyber protection of non-military government agencies, and to the best of the committee's knowledge, neither DHS nor any other part of government has been given the authority to conduct active defense on behalf of these agencies.

The primary difference between protection for government agencies and for the private sector is the fact that the actions of government agencies are subject to government control and direction within the limits of statutory law and constitutional restraint, whereas the U.S. government has exercised little influence apart from the bully pulpit today to direct or even influence the actions of much of the private sector regarding cybersecurity, a notable exception being private sector companies that are subject to strong government regulation, such as the financial sector or companies in the defense industrial base, or that provide key services to the federal government. (Whether this "hands-off" stance of government toward the private sector will continue to be the case in the future is not clear.)

This difference is perhaps most important regarding issues related to the determination of the threshold at which an active defense is appropriate. A private party setting the threshold is highly unlikely to take into account the overall national interest if it is given a free hand in determining that threshold, whereas a decision by government on the threshold is supposed to do so, at least in principle. That is, an action by a government agency is, in principle, coordinated throughout the federal government and that interagency process is responsible for ensuring that the overall national interest is indeed served by that action.

A variety of pragmatic issues are also easier to resolve when government agencies are at issue. For example, through executive order the President can direct federal agencies to share data about the scope and nature of any cyberattacks they are experiencing. Such information is necessary to determine the overall scope of any attack and thus to determine the nature of any active defense required. But most private parties are currently under no such obligation to provide such information to the federal government.²¹

²¹ Certain private parties subject to government regulation may be required to provide information under some circumstances—financial institutions, for example, are required to notify their regulatory authorities if they experience significant cyber penetrations, although this requirement is not a real-time requirement.

Decision Making and Oversight

This chapter describes decision making about and oversight of cyberattack as an instrument of U.S. national policy, focusing on issues usually associated with the Department of Defense and intelligence communities.

6.1 EXECUTIVE BRANCH

The discussion below—addressing declaratory policy, acquisition policy, and employment policy—draws from discussions of nuclear history and policy,¹ not because cyberweapons and nuclear weapons are similar (they are not), but because such discussions have highlighted the importance of several issues discussed below. That is, the committee found that nuclear history and policy are useful points of departure—framing notions and metaphorical checklists—for understanding policy regarding cyberattack but not that the conclusions that emerge from nuclear policy and history are directly applicable.

¹ Robert S. Norris, “The Difficult Discipline of Nuclear History: A Perspective,” a presentation at the Carnegie Conference on Non-Proliferation, November 7, 2005, available at http://www.carnegieendowment.org/static/npp/2005conference/presentations/Norris_Nuclear_History_Slides.pdf, and David M. Kunsman and Douglas B. Lawson, *A Primer on U.S. Strategic Nuclear Policy*, Sandia National Laboratories, Albuquerque, N.Mex., January 2001, available at http://www.nti.org/e_research/official_docs/labs/prim_us_nuc_pol.pdf.

6.1.1 Declaratory Policy

6.1.1.1 The Need for Declaratory Policy

Declaratory policy states, in very general terms, why a nation acquires certain kinds of weapons and how those weapons might be used. For example, the declaratory policy of the United States regarding nuclear weapons is stated in *The National Military Strategy*, last published in 2004:²

Nuclear capabilities [of the United States] continue to play an important role in deterrence by providing military options to deter a range of threats, including the use of WMD/E and large-scale conventional forces. Additionally, the extension of a credible nuclear deterrent to allies has been an important nonproliferation tool that has removed incentives for allies to develop and deploy nuclear forces.

By contrast, the declaratory policy of Israel regarding nuclear weapons is that it will not be the first nation to introduce nuclear weapons in the Middle East. The declaratory policy of China regarding nuclear weapons is that it will not be the first to use nuclear weapons under any circumstances. The Soviet Union once had a similar “no first use of nuclear weapons” declaratory policy, but Russia has since explicitly revoked that policy. U.S. declaratory policy has also evolved since 1945—“massive retaliation,” “flexible response,” and “escalation dominance” are some of the terms that have characterized different versions of U.S. declaratory policy regarding nuclear weapons in that period.

Declaratory policy is not necessarily linked only to the use of nuclear weapons. In 1969, the United States renounced first use of lethal or incapacitating chemical agents and weapons and unconditionally renounced all methods of biological warfare.³ In 1997, the United States ratified the Chemical Weapons Convention, which prohibits the signatories from using lethal chemical weapons under any circumstances.

Declaratory policy is directed toward adversaries as much as it is to the declaring nation itself. A declaratory policy is intended, in part, to signal to an adversary what the declaring nation’s responses might be under various circumstances. On the other hand, a declaratory policy may also be couched deliberately in somewhat ambiguous terms, leaving somewhat vague and uncertain the circumstances under which the declaring nation would use nuclear weapons. Such vagueness and uncertainty have historically been regarded by the United States as a strength rather than

² Joint Chiefs of Staff, *The National Military Strategy of the United States of America*, 2004, available at <http://www.strategicstudiesinstitute.army.mil/pdffiles/nms2004.pdf>.

³ See <http://www.state.gov/t/ac/trt/4718.htm>.

a weakness of such policies, on the grounds that uncertainty about a U.S. response is an essential part of deterring other nations from taking hostile action against its interests. By contrast, a declaratory policy that is highly explicit may be perceived as limiting a nation's options in a crisis and telegraphing its intent to some extent, thus simplifying an adversary's planning process.

Yet another related issue is whether another nation should believe a nation's declaratory policy. For example, the Soviet Union formally adopted an explicit "no-first-use" policy regarding nuclear weapons in 1982, but many military analysts gave little credence to that statement. On one hand, no immutable law mandates consistency between prior declaratory policy and subsequent action, and declaratory policy need not constrain actual practice. On the other hand, declaratory policy may influence a nation's armed forces' training and doctrine. If, for example, the declaratory policy states that a nation will not use weapon X, and its armed forces do not train to use weapon X, and its military doctrine does not contemplate the use of weapon X, that nation may well be ill-prepared to use weapon X in practice even if its leaders decide to act in violation of the stated declaratory policy.

6.1.1.2 Present Status

For the use of cyberweapons, the United States has no declaratory policy, although the DOD Information Operations Roadmap of 2003 stated that "the USG should have a declaratory policy on the use of cyberspace for offensive cyber operations." The 2006 *National Military Strategy for Cyberspace Operations* indicates that "as a war-fighting domain . . . cyberspace favors the offense . . . an opportunity to gain and maintain the initiative."⁴ This statement is the beginning of a declaratory policy, but it is incomplete.

A declaratory policy would have to answer several questions.

- For what purposes does the United States maintain a capability for cyberattack?
- Do cyberattack capabilities exist to fight wars and to engage in covert intelligence or military activity if necessary, or do they exist primarily to deter others from launching cyberattacks on the United States?
- If they exist to fight wars, are they to be used in a limited fashion?

On the basis of what is known publicly, it is possible to formulate what might be called an implied declaratory policy of the United States on cyberwarfare. (Of course, the notion of an implied declaratory policy

⁴ See <http://www.dod.mil/pubs/foi/ojcs/07-F-2105doc1.pdf>.

is itself an oxymoron—a declaratory policy that is not explicitly stated is hardly declaratory. Rather, what follows below is an example of a declaratory policy that would be consistent with what is known publicly.)

The United States acquires cyberattack capabilities as part of its overall deterrent posture, which is based on full spectrum dominance—the ability to control any situation or defeat any adversary across the range of military operations. Cyberattack capabilities provide the U.S. military and intelligence communities with additional options for action and use, and are thus intended for use just as any other weapons could be used in support of U.S. military or intelligence objectives. Cyberattack capabilities are to be fully integrated into U.S. military operations when appropriate, and distinctions between cyberattack and kinetic force are not meaningful except in an operational context. Cyberattack capabilities may be particularly useful to the United States in many conflict scenarios short of all-out war.

In addition, two other questions are often included under the rubric of declaratory policy:

- How is cyberconflict to be stopped?
- To the extent that cyberattack is part of the U.S. deterrent posture, how can its use be established as a credible threat?

In the nuclear domain, concerns have always been raised about nuclear strikes against an adversary's strategic command and control system. The issue has been that such strikes could seriously impair war termination efforts by disconnecting the political leadership of a nation from the nuclear-armed forces under its control, leaving the question of how nuclear hostilities might be terminated.

The use of large-scale cyberattacks against the communications infrastructure of an adversary might well lead to similar concerns. Such attacks could result in the effective disconnection of forces in the field from the adversary's national command authority, or sow doubt and uncertainty in an adversary's military forces about the reliability of instructions received over their communications infrastructure. Again, under such circumstances, termination of hostilities might prove problematic (and if the adversary were a nuclear-armed nation, sowing such doubt might seriously run counter to U.S. interests).

Regarding the credibility of nuclear use, the United States does much through its declaratory (and acquisition) policy to encourage the perception that there are circumstances under which the United States *might* use nuclear weapons, and it conducts large-scale military exercises involving nuclear forces in part to demonstrate to the world that it is capable

of mustering nuclear forces that could be brought to bear in any given situation.

The situation is entirely reversed with respect to cyberwarfare. U.S. policy regarding the use of cyberweapons is shrouded in secrecy, and the lack of public discussion regarding U.S. policy in this domain almost by definition does not contribute to deterrence.

Finally, the *National Military Strategy of the United States of America* of 2004 also states:⁵

The term WMD/E relates to a broad range of adversary capabilities that pose potentially devastating impacts. WMD/E includes chemical, biological, radiological, nuclear, and enhanced high explosive weapons as well as other, more asymmetrical “weapons.” They may rely more on disruptive impact than destructive kinetic effects. For example, cyber attacks on US commercial information systems or attacks against transportation networks may have a greater economic or psychological effect than a relatively small release of a lethal agent.

Coupled with the declaratory policy on nuclear weapons described earlier, this statement implies that the United States will regard certain kinds of cyberattacks against the United States as being in the same category as nuclear, biological, and chemical weapons, and thus that a nuclear response to certain kinds of cyberattack (namely, cyberattacks with devastating impacts) may be possible. It also sets the relevant scale—a cyberattack that has an impact larger than that associated with a relatively small release of a lethal agent is regarded with the same or greater seriousness.

6.1.1.3 Alternative Declaratory Policies

Simply as illustration (and not as endorsement), the following discussion incorporates and addresses hypothetical declaratory policies (or elements thereof) regarding cyberattack.

- *No large-scale cyberattacks.* Although weapons for cyberattack are valid and legitimate military weapons to be deployed and used in support of U.S. interests, the United States will unilaterally refrain from conducting against nations cyberattacks that would have the potential for causing widespread societal devastation and chaos. Accordingly, the United States will refrain from conducting cyberattacks against a nation’s electric power grids and financial

⁵ Joint Chiefs of Staff, *The National Military Strategy of the United States of America*, 2004, available at <http://www.strategicstudiesinstitute.army.mil/pdffiles/nms2004.pdf>.

systems if such attacks would have a significant potential for affecting national economies.

Such a policy would seek to delegitimize the use of large-scale cyberattacks as an instrument of national policy by any nation in much the same way that the unilateral U.S. renunciation of biological weapons contributed to stigmatizing use of such weapons by any nation. The benefit to the United States if such stigmatization occurred would be a lower likelihood that it would experience such an attack.

- *No first use of large-scale cyberattacks.* Although weapons for cyberattacks are valid and legitimate military weapons to be deployed and used in support of U.S. interests, the United States will not be the first nation in a conflict to conduct against nations cyberattacks that would have the potential of causing widespread societal devastation and chaos. Nevertheless, the United States reserves the right to conduct such attacks should it be subject to such attacks itself.

Such a policy would seek to discourage the use of large-scale cyberattacks as an instrument of national policy by any nation. However, the U.S. stance on the use of large-scale cyberattacks would be based primarily on threatening in-kind retaliation rather than setting an example. As in the previous case, the benefit to the United States if such stigmatization occurred would be a lower likelihood that it would experience such an attack.

- *No first use of cyberattacks through the Internet and other public networks.* The U.S. government will refrain from using the Internet and other public networks to conduct damaging or destructive acts, and will seek to prevent individuals and organizations within its authority from doing so, as long as other nations do the same.

Such a policy would seek to discourage the use of cyberattacks through the Internet as an instrument of national policy by any nation, presumably based on a rationale that sees the Internet as a global public utility whose benefits to the world's nations are outweighed by any temporary military advantage that might be gained through Internet-based cyberattacks. Again, the U.S. stance on the use of such cyberattacks would be based primarily on threatening in-kind retaliation rather than example-setting. The benefit to the United States would be that it (and especially its civilian sector) would be more likely to continue to enjoy the benefits of Internet connectivity.

- *National responsibility for cyberattacks.* Nations are responsible for cyberattacks that emanate from their soil, whether or not their national governments have initiated such actions. If they have not, national governments are responsible for taking actions that lead or help lead to the cessation of such actions. The United States reserves the right to take unilateral action if a nation fails to take action to respond to cyberattacks emanating from its soil.

Such a policy would codify for cyberattack a legal principle that is foundational to international law regarding neutrality, self-defense, and the laws of armed conflict (discussed further in Chapter 7)—that nations are responsible for military conduct emanating from their territories and affecting other nations. The benefit of such a policy would be to make explicit what is already U.S. policy regarding kinetic attacks.

6.1.1.4 The Relationship Between Declaratory Policy and International Agreements

Declaratory policy might also be replaced or complemented by bilateral or multilateral agreements, much as nations have sometimes agreed to certain standards of behavior for their navies on the high seas when interacting with the navies of nations also party to those agreements. This point is addressed in more detail in Chapter 10.

6.1.2 Acquisition Policy

The acquisition of capabilities is, in principle, driven by statements of need—that is, how the U.S. military (for instance) may effectively take advantage of a given capability. Much has been written about the drivers of military acquisition, and a key driver that emerges from these writings is the anticipation that an adversary has or will acquire a particular military capability to which the nation must respond quickly by itself acquiring a similar or countering capability.⁶

Acquisition policy addresses issues such as how much should be spent on weapons of various kinds, how many of what kind should be acquired on what timetable, and what the characteristics of these weapons should be. A statement of acquisition policy regarding nuclear weapons might say something like “the United States must deploy in the next 2 decades 500 land-based new ICBMs with 10 nuclear warheads apiece,

⁶ See, for example, Stephen Rosen, Chapter 7, “What Is the Enemy Building?” in *Winning the Next War: Innovation and the Modern Military*, Cornell University Press, Ithaca, N.Y., 1991.

each with a kill probability (P_k) of 90 percent against targets hardened to withstand overpressures of 2000 pounds per square inch." For a standoff munition, a statement of acquisition policy might say something like "the United States must acquire, at a rate of 1000 per year, a standoff 'fire-and-forget' munition carrying a 250-pound explosive warhead capable of being launched from a range of 30 kilometers with a Circular Error Probable of 1 meter against moving targets under all weather and battlefield conditions."

The acquisition process also requires that a weapon in acquisition be subject to an internal review prior to production to determine if use of the weapon would conflict with existing international obligations (e.g., arms control treaties or customary international standards of necessity, proportionality, and discrimination in the law of armed conflict). Not surprisingly, such review is undertaken using DOD interpretations of the law of armed conflict, which outside analysts sometimes criticize as being overly narrow. These reviews are generally not classified, but in general, they have not been made widely available.

Finally, the acquisition process requires that certain weapons undergo operational testing and evaluation before large-scale production. Operational testing and evaluation (OT/E) involves field testing under realistic combat conditions for the purpose of determining the effectiveness and suitability of a weapon for use in combat by typical military users. However, only weapons procured through a major defense acquisition program are subject to this OT/E requirement, and in particular weapons procured through a highly sensitive classified program (as designated by the secretary of defense) are exempt from this requirement.

In principle, this process also applies to the acquisition of cyberweapons, or more precisely, capabilities for cyberattack. (It would be rare that a "cyberweapon" takes the same form as a kinetic weapon, in the sense of a package that can be given to a military operator as a rifle or a fighter jet can be given. Rather, operators who launch cyberattacks are likely to have a variety of tools at their disposal for conducting an attack.) But acquiring capabilities (tools) for cyberattack differs in important ways from acquiring ordinary weapons, raising a number of issues for the acquisition process.

For example, the rapid pace of information technology change places great stress on acquisition processes for cyberattack capabilities (and for cyberdefense as well). A second important point is that the acquisition cost of software-based cyberattack tools is almost entirely borne in research and development, since they can be duplicated at near-zero incremental cost. By contrast, procurement is a major portion of the acquisition cost for kinetic weapons. Thus, a testing and evaluation (T/E) regime timed to occur after R&D is unlikely to apply to cyberweapons. The absolute

acquisition cost of cyberweapons is also likely to be significantly smaller than those of kinetic weapons, thus exempting cyberweapons from T/E regimes linked to acquisition cost.⁷

A third point is that the acquisition process presumes that it is the only way to procure weapons. But cyberattack capabilities are so inexpensive to acquire that they could be acquired through operations and maintenance (O/M) funds (and may be legal as well). For example, under the rubric of upgrading the cybersecurity posture of an installation, a system administrator might well obtain tools designed to test its computer security (that is, to support a “red team” penetration test) and acquire these tools through O/M funds. But these very same tools could provide capabilities that could be used against adversary computers.

A second way to acquire cyberattack capability is to purchase services that provide them. For example, botnets (discussed in Section 2.2.5.1.1) can be rented at relatively low cost—informed estimates vary, but are reported to be on the order of a few thousand dollars for a botnet consisting of tens of thousands of zombies for a few days. Renting a botnet may be a much more efficient method for acquiring the afforded capabilities than developing a botnet on one’s own, and indeed the Estonian minister of defense has asserted that the cyberattack on Estonia was conducted by botnets that were rented for that purpose.⁸

Of course, the rental of botnets contributes to the furtherance of a criminal enterprise, as the botnet owner/operator has broken U.S. law in assembling the botnet (presuming the owner/operator is subject to U.S. jurisdiction). An important policy question is whether it is appropriate for the United States to work with known criminals to pursue foreign policy objectives. More generally, the United States could “outsource” certain kinds of cyberattack to criminal hackers, especially if it wanted to leave no trace of such work, and incentivize such work by allowing the hackers to keep some or all of the financial resources they might encounter. Such cooperation has some precedent in U.S. history—for example, the Central Intelligence Agency sought to recruit the Mafia in 1960 to kill Fidel Castro⁹—though such instances have hardly been uncontroversial.

Related is the fact that the computers of third parties, such as innocent

⁷ For example, a major defense acquisition program is defined by statute as one estimated to require an eventual total expenditure for research, development, testing, and evaluation of more than \$300 million (in FY 1990 constant dollars) or an eventual total expenditure for procurement of more than \$1.8 billion (in FY 1990 constant dollars). Programs for acquiring cyberattack capabilities and tools are likely to cost far less than these amounts.

⁸ William Jackson, “Cyberattacks in the Present Tense, Estonian Says,” *Government Computing News*, November 28, 2007, available at http://www.gcn.com/online/vol1_no1/45476-1.html.

⁹ Glenn Kessler, “Trying to Kill Fidel Castro,” *Washington Post*, June 27, 2007, p. A06.

civilians in a nation of choice, might also be compromised in order to support a cyberattack. These computers can be configured as “weapons for cyberattack” at will by the real attacker at essentially zero cost, even though they increase his attack capabilities by orders of magnitude, and because such scenarios were never envisioned by the traditional acquisition process, it is only a matter of policy that might inhibit the United States from doing so.

Acquisition policy should also address the issue of the proper balance of resource allocation. The absolute budget sums involved in acquiring cyberattack capabilities are relatively small, as noted in Chapter 2. But serious defensive efforts are very expensive, not least for reasons of scale—the sheer volume of computer systems and networks that must be protected. Thus, acquisition policy necessarily affects the balance between conventional military assets and cyber military assets and procedures on the defensive side. Given the dependence of today’s military forces on information technologies, some analysts have argued that present-day acquisition policies do not pay sufficient attention to cybersecurity and defensive operations.

The above discussion of acquisition policy relates primarily to the defense community. But the intelligence community must also acquire various capabilities to support its intelligence collection and covert action missions. Of particular significance for acquisition policy is that a tool to collect intelligence information from an adversary computer system or network can—at little additional cost—be modified to include certain attack capabilities, as described in Section 2.6. Indeed, the cost of doing so is likely to be so low that in the most usual cases, acquisition managers would probably equip a collection tool with such capabilities (or provide it with the ability to be modified on-the-fly in actual use to have such capabilities) as a matter of routine practice.

6.1.3 Employment Policy

Employment policy specifies how weapons can be used, what goals would be served by such use, and who may give the orders to use them. Such policy has a major influence on how forces train (e.g., by driving the development and use of appropriate training scenarios).

One key question of employment policy relates to the necessary command and control arrangements. For example, although U.S. doctrine once did not differentiate between nuclear and non-nuclear weapons,¹⁰

¹⁰ In 1954, President Eisenhower was asked at a press conference (March 16, 1954) whether small atomic weapons would be used if war broke out in the Far East. He said, “Yes, of course they would be used. In any combat where these things can be used on strictly mili-

this is most surely not the case today. Nuclear weapons are universally regarded as worthy of special attention, policies, and procedures, and their use is tightly controlled and highly centralized—more so than any other weapon in the U.S. arsenal. Whether similar arrangements will be made for cyberweapons in the future remains to be seen, although the discussion in Chapter 3 suggests that the command and control arrangements of today are not as centralized.

A second key question of employment policy is the targets of such weapons. Some targets are off-limits by virtue of the LOAC and other relevant international law. But the propriety of attacking other kinds of targets is often determined by doctrine and views of the adversary.

For example, in the nuclear strategy of the Cold War, considerable debate arose about the propriety of targeting adversary nuclear forces. Advocates of prompt hard-target kill capabilities (that would use a ballistic missile against a hardened adversary missile silo) argued that the adversary (generally the leaders of the Soviet Union) placed great value on their instruments of national power, such as their nuclear forces, and that placing such instruments at risk would help to deter actions that worked against the interests of the United States. Opponents of such targeting argued that threatening to destroy such targets only increased the likelihood that the adversary would launch its missiles on warning of attack, thus making accidental launch more likely.

Given that there are no cyber equivalents of hardened missile silos that constitute an adversary's retaliatory forces, no credible threat of annihilation, and no equivalent of launch on warning for cyber forces, nuclear strategy does not provide guidance for cyber targeting. What targets might or might not be appropriate for cyberattack and under what circumstances would this be so? From what can be determined from public statements, the DOD believes that cyberattack has military utility, and thus the use of cyberattack is subject to constraints imposed by the law of armed conflict.

At the same time and apart from the need to comply with the LOAC, good reasons may exist for eschewing certain kinds of cyberattack against certain kinds of target for reasons other than those related to operational efficacy. For example, cyberwarfare provides tools that can be focused directly on messaging and influencing the leadership of an adversary

tary targets and for strictly military purposes, I see no reason why they shouldn't be used just exactly as you would use a bullet or anything else." (See Eisenhower National Historic Site, National Park Service, at <http://www.nps.gov/archive/eise/quotes2.htm>.) Indeed, in 1953, the U.S. National Security Council noted that "in the event of hostilities, the United States will consider nuclear weapons to be as available for use as other munitions." (U.S. National Security Council (NSC), "Basic National Security Policy," NSC Memorandum 162/2, October 30, 1953, available at <http://www.fas.org/irp/offdocs/nsc-hst/nsc-162-2.pdf>.)

state. Message-based influence might help to persuade the leadership to make decisions helpful to U.S. national interests, such as terminating hostilities or refraining from using weapons of mass destruction. But at the same time, it may be undesirable to conduct destructive or disruptive attacks on the command and control systems that connect the adversary's national command authority to forces in the field.

Disconnecting an adversary's forces from their leadership may result in serious dysfunction, uncoordinated action, and psychological impact on the adversary such as fear and poor morale. Such positive effects must be balanced against possible negative effects, such as the inability of the adversary's leadership to direct its forces to surrender or to stand down. In addition, if forces in the field lose confidence in the authoritative nature of commands from their national command authority, they may resort to following standing orders issued before the conflict began—and such orders may well instruct these forces to act in more destructive ways than they otherwise would. These considerations are particularly important if the adversary has nuclear weapons and if the cyberattack cannot differentiate between command and control systems for the adversary's conventional and nuclear forces.

Other possible targets to be avoided may include those that could have significantly damaging effects on large numbers of non-combatants. Entirely apart from the moral and ethical issues raised by such attacks, conducting such attacks against a nation with a declared policy of responding to such attacks with nuclear weapons arguably increases the likelihood that such weapons would be used. Targets in this category might include national financial systems and electric power grids.

Cyberattacks may be a preferred method for targeting infrastructure under some circumstances. The United States may wish to conduct operations related to war recovery and stabilization in the aftermath of a conflict, and thus wish to preserve infrastructure as an important element in war recovery—the U.S. intent in Operation Iraqi Freedom (the Second Gulf War) in 2003 was to occupy Baghdad for some period of time thereafter and to enable Iraq to function as a sovereign nation. In its targeting of Iraqi infrastructure, the United States had to consider the possibility that destroying parts of it (e.g., the electric power grid) might impede war recovery efforts after the conflict. If cyberattacks made it possible to attack infrastructure in such a way that it was rendered non-functional for the duration of a conflict but could be easily restored to normal operation after the conflict was terminated, attack planners would have considerable incentives to prefer such attacks over more destructive ones.

A second issue relates to options for strategic use. As with nuclear weapons, the availability of preplanned options for cyberattack varying in scale, scope, and timing would increase flexibility and the ability to

respond promptly to various strategic contingencies. A number of important questions arise in this context—the large amount of intelligence information likely to be needed for such options, the timeliness of information collected to support preplanned options, and indeed the actual value of prompt cyber response under various circumstances.

A third important issue is ensuring that cyberattack activities are sufficiently visible to higher authorities, including the political leadership. It is an unfortunate reality that during times of crisis, military actions that would normally be regarded as routine or “small” can lead to misperceptions of strategic significance. For example, routine air reconnaissance undertaken during times of crisis can be interpreted as a prelude to attack. In a cyberattack context, analogs could include the routine gathering of intelligence that is needed to support a cyberattack (e.g., port scans of Zendian systems) or the self-defense neutralization of an active cyberattack threat from a Zendian patriotic hacker under standing rules of engagement. The possibility is very real that Zendian authorities might perceive such activities as aggressive actions associated with a planned and deliberate cyberattack by the United States.

Keeping the political leadership informed of such activities is a problem even when considering traditional military operations. But because the resources and assets needed to conduct cyberattacks are small by comparison and the potential impact still large, it may be more difficult for higher authorities to stay informed about activities related to cyberattack.

Finally, the United States has a long-standing policy not to use cyberattack or cyberexploitation to obtain economic advantage for private companies (as noted in Section 4.1.2). However, the economic domain is one in which the operational policies of adversaries are markedly different from those of the United States. That is, adversaries of the United States are widely believed to conduct cyber-espionage for economic advantage—stealing trade secrets and other information that might help them to gain competitive advantage in the world marketplace and/or over U.S. firms. As noted in Section 2.6.2, the intelligence services of at least one major nation-state were explicitly tasked with gathering intelligence for its potential economic benefits. This asymmetry between U.S. and foreign policies regarding cyberexploitation is notable.

The committee also observes that national policy makers frequently refer to a major and significant cyberthreat against the United States emanating from many actors, including major nation-states. The result in recent years has been an upsurge of concern about the disadvantaged position of the United States in the domain of cyberconflict, and is most recently reflected in the still largely classified Comprehensive National Cybersecurity Initiative resulting from the National Security Presiden-

tial Directive 54/Homeland Security Presidential Directive 23 of January 2008.¹¹

On the other hand, the committee's work has underscored many of the uncertainties that underlie any serious attempt by the United States to use cyberattack as an instrument of national policy. Moreover, military planners often engage in worst-case planning, which assumes that more things will go right for an adversary than for oneself. Thus, attack planners emphasize the uncertainties of an attack and assume that the defense will be maximally prepared and lucky. Defensive planners emphasize the uncertainties of defense and assume that the attacker will be maximally prepared and lucky.

In short, the committee sees a marked asymmetry in the U.S. perception of cyberattack—"they" (the adversary) are using cyberattack means effectively against us (the United States), but it would be difficult (though not impossible) for us to use such means effectively against them.

The question thus arises, What might be responsible for this perception? One factor is the conflation of cyberattack and cyberexploitation in the public discourse (see Box 1.4 in Chapter 1). As noted by General Kevin Chilton, commander of the U.S. Strategic Command, many of the incidents that are billed as cyberattacks are, more accurately, just old-fashioned espionage—people looking for information who don't necessarily represent military threats.¹² Thus, if the public discourse uses the term "cyberattack" (what this discussion calls cyberattack-AUIPD, for "cyberattack as used in public discourse," to distinguish usages) to include cyberexploitation, then the balance is between adversary cyberattacks-AUIPD (which would include what this report terms "cyberattack" [note absence of a tag] and which are largely espionage conducted for economic benefit) and U.S. "cyberattacks-AUIPD" (which by policy do not involve either cyberattack or cyberexploitation conducted for economic benefit), and in such a balance, adversary cyberattacks-AUIPD will obviously seem to be much more effective than those of the United States.

A third important factor contributing to this perception is the fact

¹¹ Public reports indicate that this initiative has 12 components intended to reduce to 100 or fewer the number of connections from federal agencies to external computer networks, and to make other improvements in intrusion detection, intrusion prevention, research and development, situational awareness, cyber counterintelligence, classified network security, cyber education and training, implementation of information security technologies, deterrence strategies, global supply chain security, and public/private collaboration. The cost of this initiative has been estimated at \$40 billion. See, for example, Jill R. Aitoro, "National Cyber Security Initiative Will Have a Dozen Parts," *Nextgov*, August 1, 2008, available at http://www.nextgov.com/nextgov/ng_20080801_9053.php.

¹² Wyatt Kash, "Cyber Chief Argues for New Approaches," *Government Computer News*, August 22, 2008, available at <http://gcn.com/articles/2008/08/22/cyber-chief-argues-for-new-approaches.aspx>.

that as noted in earlier chapters, the United States provides only limited assistance to the private sector when it comes under cyberattack and restricts the ability of the private sector to engage in self-help activities (as discussed in Section 5.2), and it refrains from sharing intelligence information that would benefit individual private sector companies (as discussed in Section 4.1). Some other nations do not practice such restraint. The committee speculates that this asymmetry in policy may account for at least some of the perception of asymmetric advantage derived by others.

If these observations are accurate, what—if anything—can be done about it?

Regarding the conflation of cyberattack and cyberexploitation in public discourse, there is no remedy except to insist that a user of the term “cyberattack” make clear what is included under the rubric of the term he or she is using. If the many foreign cyberexploitation efforts were not described as “cyberattack,” the level of tension over cyberattack would be knocked down to a considerable degree.

The case for the current U.S. policy regarding eschewing the use of U.S. intelligence agencies for the benefit of private firms is largely based on the desire of the United States to uphold a robust legal regime for the protection of intellectual property and for a level playing field to enable competitors from different countries to make their best business cases on their merits. If this policy position is to be revised, it seems that two of the most prominent possibilities are that (1) intelligence gathering for economic purposes ceases for all nations, or (2) the United States uses its intelligence-gathering capabilities (including cyberexploitation) for economic purposes. Under traditional international law, espionage—for whatever purpose—is not banned, and thus the first possibility suggests a need to revise the current international legal regime with respect to the propriety of state-sponsored economic espionage. The second possibility raises the prospect that current restraints on U.S. policy regarding intelligence collection for the benefit of private firms might be relaxed.

Both of these possibilities would be controversial, and the committee takes no stand on them, except to note some of the problems associated with each of them. The first—a change in the international legal regime to prohibit espionage—would require a consensus among the major nations of the world, and such a consensus is not likely. The second—a unilateral change in U.S. policy—does not require an international consensus, but has many other difficulties. For example, the U.S. government would have to decide which private firms should benefit from the government’s activities, and even what entities should count as a “U.S. firm.” U.S. government at the state and local level might well find that the prospect of U.S. intelligence agencies being used to help private firms would not sit well with foreign companies that they were trying to persuade to relocate

to the United States. And it might well undercut the basis on which the United States could object to other nations conducting such activities for the benefit of their own domestic industries and lead to a “Wild West” environment in which anything goes.

After all is said and done, it may turn out that the most desirable (least undesirable) option for the United States is to learn to live with the current asymmetry. But if that is indeed the case, it should reflect a deliberate and considered assessment of the pros and cons of various options that in the committee’s view has not yet been engaged.

6.1.4 Operational Oversight

Operations translate employment policy into reality. In practice, the U.S. armed forces operate on a worldwide basis and have many ongoing operations at any given time. For example, they constantly gather intelligence and reconnaissance information. Some of those operations are sensitive, in that they might be seen as provocative or otherwise inappropriate.

Thus, the U.S. government has established a variety of mechanisms intended to ensure that such operations are properly overseen. For example, the U.S. government sometimes specifies criteria in advance that define certain sensitive military missions, and then requires that all such missions be brought to the attention of senior decision makers (e.g., the National Security Council staff). In rare cases, a mission must be approved individually; more typically, generic authority is granted for a set of missions that might be carried out over a period of many months (for example). The findings and notification process for covert action is another mechanism for keeping the executive and legislative branches properly informed. From time to time these mechanisms are unsuccessful in informing senior decision makers, and it is often because the individual ordering the execution of that mission did not believe that such an order required consultation with higher authority.

In a cyberattack context, oversight issues arise at two stages—at the actual launch of a cyberattack and in activities designed for intelligence preparation of the battlefield to support a cyberattack.

6.1.4.1 Launching a Cyberattack

Another important operational issue involves delegation of authority to launch a cyberattack as part of an active defense of U.S. computer systems and networks. As noted in Chapter 3, the U.S. Strategic Command has authority to conduct such attacks for active defense under a limited set of circumstances. But it is not known how far down the chain of command such authority has been delegated.

The most extreme form of delegation would call for an entirely automated active defense—and indeed the U.S. Air Force has issued a call for proposals to develop a “cyber control system” that “will enable active defense operations [involving] automated responses (based on predefined Rules of Engagement) . . . , in response to network intrusions/attacks.”¹³ Automated responses are regarded as being militarily necessary when there is insufficient time for humans to make decisions about the nature of a response and any given situation may present insufficient time because of the fleeting nature of the opportunity to strike back or because of the harm that rapidly accrues if the attack is not stopped (though consideration of other factors such as appropriate rules of engagement may prevent such weapons from being deployed in any given situation). Both of these factors could characterize certain kinds of cyberattacks on certain targets in the United States.

On the other hand, the risks of error or inadvertent escalation are generally regarded as greatest when humans are not in the decision-making loop. Despite periodic calls for the nuclear command and control system to be automated so as to ensure that retaliation would take place in the event of a Soviet nuclear attack, the United States has always relied on humans (the President and the National Command Authority) to make the ultimate decision to release U.S. strategic forces. (Even so, many have criticized these arrangements as *pro forma*, arguing that in practice they are not much better than an automated launch decision, because they give the NCA too little time to evaluate the information available about the alleged incoming attack.)

An assessment of the wisdom of an automated response to a cyberattack depends on several factors, including the likelihood that adequate and correct information will be available in a short period of time to develop an access path back to the attacker, the likely consequences of a cyberattack response, and the possible consequences of a misdirected or inappropriately launched counterattack. In the case of nuclear command and control, these factors—primarily the last—indicate that an automated response would be foolish and foolhardy.

6.1.4.2 Conducting Intelligence Preparation of the Battlefield to Support a Cyberattack

In principle, conducting intelligence preparation of the battlefield (IPB) to support a cyberattack is not different from conducting other non-destructive cyberexploitation missions. For example, U.S. electronic

¹³ United Press International, “Air Force Seeks Automated Cyber-response,” Jan. 2, 2008, at 4:55 p.m.

reconnaissance airplanes often fly missions near the border of another nation in order to “light up” that nation’s air defense radars. By monitoring those radar emissions, they collect information on the waveforms and positions of a potential adversary’s radar systems; such information could be useful in the event that an air strike might be launched against that nation.

On the other hand, that nation might well regard those reconnaissance flights as provocative. The airplane it is monitoring just outside its airspace could be armed, and the plane’s presence there could indicate hostile intent. The essential problem is that the boundaries of its national airspace provide almost no time for its air defense forces to react should the airplane turn out to have immediate hostile intent. Even if it is known to be unarmed, it is most likely to be a reconnaissance airplane collecting information that could be useful in the event that an air strike was launched against that nation. If these reconnaissance flights were taking place during a period of peacetime tension with the United States, it is easy to see how they might further exacerbate those tensions.

Missions of this kind fall squarely into the category of those that must be reported to senior policy makers. The IPB mission for a destructive cyberattack falls into the same category. In order to gather the necessary intelligence, an adversary’s network must be mapped to establish topology (which nodes are connected to which other nodes). Ports are “pinged” to determine what services are (perhaps inadvertently) left open to an outside intruder, physical access points are located and mapped, operating system and application vulnerabilities are identified, sympathizers with important access privileges are cultivated, and so on.

However, there are at least three important differences between IPB for cyberattack and other kinds of intelligence collection. First, a U.S. government effort to conduct IPB for many kinds of cyberattack will be taking place against a background of other activities (e.g., probes and pings) that are not being conducted by the U.S. government. Second, network connectivity may be such that “limited” intelligence probes and other investigations of a potential adversary’s networks will inadvertently reach very sensitive areas. Third, the dividing line between a tool intended to collect information on an adversary’s systems and a weapon intended to destroy parts of those systems may be very unclear indeed.

The first factor above may reduce the sensitivity of the nation being probed—and indeed, the U.S. IPB effort is likely to be undertaken in a way that does not reveal its origin. But the second two factors may increase sensitivity, and possibly lead to entirely unanticipated reactions on the part of the adversary.

6.2 LEGISLATIVE BRANCH

The legislative branch has two basic roles regarding government operations—budget and oversight. In addition, the Constitution gives the legislative branch the sole authority to declare war.

6.2.1 Warmaking Powers

Article I, Section 8 of the U.S. Constitution authorizes the Congress to “declare war” and gives Congress numerous powers over the military, including the powers to “raise and support armies,” to “provide and maintain a navy,” and to “make rules for the government and regulation of the land and naval forces.” Article II, Section 2 gives the President the “executive power” and provides that he “shall be commander in chief of the Army and Navy of the United States.”

At the time the Constitution was written, the primary purpose of national armed forces was to fight wars, and these provisions were intended to give Congress primary responsibility for the decision to initiate war, and to give the President the primary responsibility for the conduct of war.¹⁴ Over time, as the international powers and responsibilities of the United States have grown, and as the standing U.S. armed forces have grown, the President has asserted more and more authority to initiate armed conflicts in the absence of authorization from Congress. Moreover, it has been argued that the notion of declaring war as a prelude to armed combat is simply irrelevant in the modern world.

Self-defense is the least controversial basis for the president to direct the armed forces to engage in combat. Madison said at the Convention that the “declare war” clause left to the President the power to “repel sudden attacks” without congressional authorization.¹⁵ The Supreme Court upheld Lincoln’s authority to act against the confederacy in the absence of congressional authorization.¹⁶ President Clinton invoked self-defense in justifying the 1993 cruise missile strikes on Iraq in response to the attempted assassination of President George H.W. Bush.¹⁷

For some of the instances not involving self-defense in which U.S. armed forces have been deployed and used, presidents have sought and

¹⁴ See, e.g., Abraham D. Sofaer, *War, Foreign Affairs and Constitutional Power: The Origins*, Ballinger Publishing, Cambridge, Mass., 1976.

¹⁵ The Records of the Federal Convention of 1787, at 318 (1911), Max Farrand, ed., rev. edition, 1966.

¹⁶ See *Prize Cases*, 67 U.S. 635 (1863) (“If a war be made by invasion of a foreign nation, the President is not only authorized but bound to resist force by force”).

¹⁷ See “Letter to Congressional Leaders on the Strike on Iraqi Intelligence Headquarters,” Pub. Papers of William J. Clinton 940, 1993.

received explicit congressional authorization, although they have always claimed that their authority as commanders-in-chief was sufficient to take such actions and that in essence seeking congressional authorization was a courtesy extended to the legislative body. But matters are more complicated and controversial when the President acts without invoking self-defense and also without congressional authorization.

The President has acted in such a manner in many circumstances in U.S. history, most notably in Korea and Kosovo, but also in dozens of other smaller-scale conflicts. Presidents have asserted this authority, Congress often complains and opposes it, and the Supreme Court has not squarely addressed it.

To address such cases, Congress passed the War Powers Resolution (WPR) in 1973 (PL 93-148). Passed over then-President Nixon's veto, the WPR requires the President to report to Congress in 48 hours "in any case in which United States Armed Forces are introduced (1) into hostilities or into situations where imminent involvement in hostilities is clearly indicated by the circumstances; (2) into the territory, airspace or waters of a foreign nation, while equipped for combat, except for deployments which relate solely to supply, replacement, repair, or training of such forces; or (3) in numbers which substantially enlarge United States Armed Forces equipped for combat [who are] already located in a foreign nation," and requires the President to "terminate any such use of armed forces" within 60 days (subject to a one-time 30-day extension).

The tensions between the executive and legislative branches of government over war-making authority are palpable. Many analysts believe that the intent of the Founding Fathers was to grant the Congress a substantial decision-making role in the use of U.S. armed forces, and if modern conflict has rendered obsolete the notion of a "declaration of war," mechanisms must still be found to ensure that Congress continues to play a meaningful role in this regard. Others acknowledge the obsolete nature of declarations of war, but conclude that executive branch authority can and should fill the resulting lacunae.

This report does not seek to resolve this controversy, but observes that notions of cyberconflict and cyberattack will inevitably cause more confusion and result in less clarity. Consider, for example, the meaning of the term "hostilities" in the War Powers Resolution. At the time the resolution was crafted, cyberattack was not a concept that had entered the vocabulary of most military analysts. In the context of the resolution, hostilities refer to U.S. land, air, and naval units engaging in combat. The resolution also refers to the foreign deployments of combat-equipped U.S. forces.

To the extent that the War Powers Resolution was intended to be a reassertion of congressional authority in warmaking, it is very poorly suited to U.S. forces that engage in cyber combat or launch cyberattacks.

What conditions would define “hostilities” when military cyberattacks can be launched against adversary computers or networks? What counts as “deployments” of forces capable of cyberattack into foreign territory? It is thus an open question whether a cyberattack launched by the United States would constitute the introduction of armed forces in another country within the meaning of the resolution.

When it comes to sorting out normative and practical issues concerning congressional and presidential prerogatives, cyberwarfare poses issues even more difficult for interpreting the War Powers Resolution than the already-difficult issues associated with traditional kinetic conflict.

6.2.2 Budget

In the preceding section, the relative invisibility of cyberattack activities is mentioned as a problem for higher authority. Cyberattack capabilities are also not particularly visible to the legislative branch. In part, the veil of secrecy around cyberattack makes it more invisible than if the subject were not classified. But just as important is the fact that the funding for the development and deployment of cyberattack capabilities is both minuscule and deliberately obscured in unclassified budget justifications.

For example, in the FY 2008 DOD budget request, one request for the “demonstration of offensive cyber operations technologies allowing attack and exploitation of adversary information systems” by the Air Force is contained in a program element component of \$8.012 million; the program element is entitled “Advanced Technology Development,” and the component “Battlespace Information Exchange.”¹⁸ A second request for developing cyber operations technologies is contained in a program element of \$11.85 million for FY 2008; this program element is entitled “Applied Research on Command, Control, and Communications.”¹⁹

A reasonable observation is that development and demonstration of cyberattack capabilities are distributed over multiple program elements,

¹⁸ See <http://www.dtic.mil/descriptivesum/Y2008/AirForce/0603789F.pdf>.

¹⁹ In FY 2008, one component of this program element (“communications technology”) called for activities to “initiate development of access techniques allowing “cyber paths” to protected adversary information systems through a multiplicity of attack vectors; initiate development of stealth and persistence technologies enabling continued operation within the adversary information network; initiate programs to provide the capability to exfiltrate any and all types of information from compromised information systems enabling cyber intelligence gathering to achieve cyber awareness and understanding; initiate technology programs to deliver D5 (deny, degrade, destroy, disrupt, and deceive) effects to the adversary information systems enabling integrated and synchronized cyber and traditional kinetic operations.” See <http://www.dtic.mil/descriptivesum/Y2008/AirForce/0602702F.pdf>.

each of which is relatively small in financial terms. Budget oversight is thus difficult to execute, even though it is intimately related to acquisition policy. In addition, the ability to increase certain attack capabilities “for free” (e.g., through the use of botnets and automated production functions) negates to a considerable extent the ability of the legislative branch to use budget totals for restraining or limiting U.S. military capabilities.

A low budget profile supports low visibility. Proponents of a given capability would prefer low visibility for programs supporting that capability, especially if the capability were controversial in nature. (Low visibility can also be achieved in other ways, such as by designating a program as “special access.”)

6.2.3 Oversight (and Notification)

In addition to budgetary oversight, the legislative branch also provides operational oversight of government programs. For example, the executive branch is required by law (50 U.S.C. 413(a)(1)) to keep the congressional intelligence committees “fully and currently informed” of all U.S. intelligence activities, including any “significant anticipated intelligence activity.”²⁰ Both intelligence gathering and covert action are included under this rubric, and thus cyberexploitation and covert action cyberattacks would have to be reported to these committees. These reporting requirements are subject to a number of exceptions pertaining to sensitivity and possible compromise of intelligence sources and methods, or to the execution of an operation under extraordinary circumstances.

Certain DOD operations have also been subject to a notification requirement. Section 1208 of the FY 2005 Defense Authorization Act gave the secretary of defense the authority to expend up to \$25 million in any fiscal year to “provide support to foreign forces, irregular forces, groups, or individuals engaged in supporting or facilitating ongoing military operations by United States special operations forces to combat terrorism.” In the event that these funds were used, the secretary of defense was required to notify the congressional defense committees expeditiously and in writing, and in any event in not less than 48 hours, of the use of such authority with respect to that operation.

Yet another precedent for notification in support of oversight is the requirement for the attorney general to report annually to Congress and the Administrative Office of the United States Courts indicating the total

²⁰ A discussion of this requirement can be found in Alfred Cumming, *Statutory Procedures Under Which Congress Is to Be Informed of U.S. Intelligence Activities, Including Covert Actions*, Congressional Research Service memo, January 18, 2006, available at <http://www.fas.org/sgp/crs/intel/m011806.pdf>.

number of applications made for orders and extensions of orders approving electronic surveillance under the Foreign Intelligence Surveillance Act, and the total number of such orders and extensions either granted, modified, or denied.

To the best of the committee's knowledge, no information on the scope, nature, or frequency of cyberattacks conducted by the United States has been made regularly or systematically available to the U.S. Congress on either a classified or an unclassified basis.

Part III

Intellectual Tools for Understanding and Thinking About Cyberattack

Cyberattack is a complex topic that is often oversimplified and/or demonized in popular accounts. Part III is intended to provide the reader with a variety of intellectual tools useful for thinking about cyberattack in its many forms and permutations.

Chapter 7 provides a primer on legal issues relevant to cyberattack, mostly related to international law. These issues arise because of the U.S. commitment to act in compliance with the international law of armed conflict and its treaty obligations (including that of the Charter of the United Nations) and because cyberattack presents various challenges in interpreting laws written and established to govern traditional kinetic warfare between nation-states.

Chapter 8 examines a number of previously studied areas for possible lessons relevant to cyberattack. These include nuclear weapons and warfare, space as a domain of conflict, biological weapons, and non-lethal weapons. Chapter 8 will demonstrate that all of these areas have within them some relevant study or precedent, but that none of them carry over fully to cyberattack.

Chapter 9 presents some speculations on the dynamics of cyberconflict as it might involve the United States as a major player. Not much is known today about how such cyberconflict might start, and even less about how it would evolve over time. The best that can be done is to reason from analogy in a largely preliminary fashion, and that is the role of Chapter 9.

Chapter 10 describes a variety of alternative futures. The current stance of the United States toward cyberattack is one that puts no constraints on its use apart from those imposed by the law of armed conflict and related customary international law. But other stances are possible, and from time to time proposals emerge that, if adopted, would constrain activities related to cyberattack undertaken by all nations, including the United States. Chapter 10 explores some of these proposals in notional form but does not take a stand one way or another on their inherent desirability.

Legal and Ethical Perspectives on Cyberattack

7.1 THE BASIC FRAMEWORK

In the context of this chapter, international law refers to treaties (written agreements among states governed by international law) and customary international law (general and consistent practices of states followed from a sense of legal obligation). Domestic law refers to the Constitution of the United States, federal statutes, and self-executing treaties¹ and can constrain the actions of government and of private individuals.

This chapter focuses on the implications of existing international and domestic law as well as relevant ethical regimes for the use of cyberattack by the United States. (It is thus not intended to address legal issues that arise mostly in the context of the United States defending against cyberattack.) Compared to kinetic weapons, weapons for cyberattack are a relatively recent addition to the arsenals that nations and other parties can command as they engage in conflict with one another. Thus, the availability of cyberattack weapons for use by national governments naturally raises questions about the extent to which existing legal and ethical perspectives on war and conflict and international relations—which affect

¹ In 2008, the Supreme Court explained that a self-executing treaty is one that “operates of itself without the aid of any legislative provision,” and added that a treaty is “not domestic law unless Congress has either enacted implementing statutes or the treaty itself conveys an intention that it be ‘self-executing’ and is ratified on these terms.” See *Medellin v. Texas*, 128 S.Ct. 1346, 1356 (2008) (citations and internal quotations omitted).

considerations of how and when such weapons might be used—could require reinterpretation or revision.

Some analysts have responded to these questions in the negative, arguing that cyberweapons are no different than any other weapons and thus that no new legal or ethical analysis is needed to understand their proper use.² Others have taken the opposite view, arguing that cyberweapons are so different from kinetic weapons that new legal regimes are needed to govern their use.³ Further, some argue that it is much easier to place substantive constraints on new military technologies before they have been integrated into the doctrine and structure of a nation's armed forces. And still others have taken the view that although cyberweapons do raise some new issues, the basic principles underlying existing legal and ethical regimes continue to be valid even though analytical work is needed to understand how these principles do/should apply to cyberweapons.

As is indicated below in this chapter, the committee's perspective is most similar to the last one articulated above. Furthermore, the committee observes that in no small measure, the range of opinions and conclusions about the need for new regimes comes from the fact that as indicated in Chapter 2, the notion of cyberattack spans an enormous range of scale, impact, and complexity. Some specification of a cyberattack's range, scope, and purpose must be presented if analytical clarity is to be achieved.

This chapter does not attempt to provide a comprehensive normative analysis. Instead, it reviews the current international and domestic legal regimes, and suggests where existing regimes may be inadequate or ambiguous when the use of cyberweapons is contemplated. In addition, it explores issues that cyberattack may raise outside the realm of the relevant legal regimes. In all instances, the emphasis is on raising questions, exploring ambiguities, and stimulating thought.

Although this report takes a Western perspective on ethics and human rights, the committee acknowledges that these views are not universal. That is, other religious and ethnic cultures have other ethical and human rights traditions and practices that overlap only partially with those of the United States or the West, and their ethical and human rights traditions may lead nations associated with these cultures to take a different perspective on ethical, human rights, and legal issues regarding cyberattack. Perhaps most importantly, other nations may take a more expansive or a

² This point of view was expressed in presentations to the committee by the USAF Cyberspace Task Force (briefing of LTC Forrest Hare, January 27, 2007).

³ See, for example, Christopher C. Joyner and Catherine Lotrionte, "Information Warfare as International Coercion: Elements of a Legal Framework," *European Journal of International Law* 12(5):825-865, 2001.

more restricted view of how the law of armed conflict constrains activities related to cyberattack.

Finally, it should be noted that legal considerations are only one set of factors that decision makers must take into account in deciding how to proceed in any given instance. There will no doubt be many circumstances in which the United States (or any other nation) would have a *legal* right to undertake a certain action, but might choose not to do so because that action would not be politically supportable or would be regarded as unproductive, unethical, or even harmful.

7.2 INTERNATIONAL LAW

International obligations flow from two sources: treaties (in this context, the Charter of the United Nations, the Hague and Geneva Conventions with their associated protocols, and the Cybercrime Convention) and customary international law. Defined as the customary practices of nations that are followed from a sense of legal obligation, customary international law has the same force under international law as a treaty.

Provisions of international law are sometimes enacted into national laws that are enforceable by domestic institutions (such as the President and courts). For example, Title 18, Section 2441 of the U.S. Code criminalizes the commission of war crimes and defines war crimes as acts that constitute grave breaches of the Geneva or Hague Conventions. Such laws impose penalties on individuals who violate the relevant provisions of international law.

When nations violate international law, the recourse mechanisms available are far less robust than in domestic law. For example, the International Court of Justice has held specific nations in violation of international law from time to time, but it lacks a coercive mechanism to penalize nations for such violations. In principle, the UN Security Council can call for coercive military action that forces a violator to comply with its resolutions, but the viability of such options in practice is subject to considerable debate.

7.2.1 The Law of Armed Conflict

To understand the legal context surrounding cyberattack as an instrument that one nation might deploy and use against another, it is helpful to start with existing law—that is, the international law of armed conflict (LOAC).

Today's international law of armed conflict generally reflects two

central ethical principles.⁴ First, a state that uses force or violence against another state must have “good” reasons for doing so, and indeed, throughout most of history, states that have initiated violence against other states have sought to justify their behavior. Second, even if violent conflict between nations is inevitable from time to time, unnecessary human suffering should be minimized.

LOAC addresses two separate questions. First, when is it legal for a nation to use force against another nation? This body of law is known as *jus ad bellum*. Second, what are the rules that govern the behavior of combatants who are engaged in armed conflict? Known as *jus in bello*, this body of law is separate and distinct from *jus ad bellum*.

7.2.1.1 *Jus ad Bellum*

Jus ad bellum is governed by the UN Charter, interpretations of the UN Charter, and some customary international law that has developed in connection with and sometimes prior to the UN Charter.

Article 2(4) of the UN Charter prohibits every nation from using “the threat or use of force against the territorial integrity or political independence of any state, or in any other manner inconsistent with the Purposes of the United Nations.” Nations appear to agree that a variety of unfriendly actions, including unfavorable trade decisions, space-based surveillance, boycotts, severance of diplomatic relations, denial of communications, espionage, economic competition or sanctions, and economic and political coercion, do not rise to the threshold of a “use of force,” regardless of the scale of their effects. As for the “threats of force” prohibited by Article 2(4), Professor Thomas Wingfield of the U.S. Army Command and General Staff College testified to the committee that such threats might plausibly include verbal threats, initial troop movements, initial movement of ballistic missiles, massing of troops on a border, use of fire control radars, and interference with early warning or command and control systems.

The UN Charter also contains two exceptions to this prohibition on the use of force. First, Articles 39 and 42 permit the Security Council to authorize uses of force in response to “any threat to the peace, breach of the peace, or act of aggression” in order “to maintain or restore international peace and security.”

⁴ The law of armed conflict is also sometimes known as international humanitarian law. A number of legal scholars, though not all by any means, view international humanitarian law as including human rights law, and thus argue that the law of armed conflict also includes human rights law. For purposes of this chapter and this report, the law of armed conflict does not include human rights law.

Second, Article 51 provides as follows: “Nothing in the present Charter shall impair the inherent right of individual or collective self-defense if an armed attack occurs against a Member of the United Nations, until the Security Council has taken measures necessary to maintain international peace and security.” The self-defense contemplated by Article 51 does not require Security Council authorization. Professor Wingfield argued that armed attack would include declared war, de facto hostilities, occupation of territory, a blockade, the destruction of electronic warfare or command and control systems, or the use of armed force against territory, military forces, or civilians abroad. In addition, there is debate over whether the right of self-defense is limited by Article 51, or whether Article 51 simply recognizes a continuation of the preexisting (“inherent”) right of self-defense. Box 7.1 elaborates on notions of self-defense and self-help.

An important aspect of the interpretation of Article 51 involves the question of imminent attack. It is widely accepted that a nation facing unambiguous imminent attack is also entitled to invoke its inherent right of self-defense without having to wait for the blow to fall. (Self-defense undertaken under threat of imminent attack is generally called “anticipatory self-defense.”) For example, *Oppenheim’s International Law: Ninth Edition* states that:⁵

The development of the law, particularly in the light of more recent state practice, . . . suggests that action, even if it involves the use of armed force and the violation of another state’s territory, can be justified as self-defence under international law where:

- a) an armed attack is launched, or is immediately threatened, against a state’s territory or forces (and probably its nationals);
- b) there is an urgent necessity for defensive action against that attack;
- c) there is no practicable alternative to action in self-defence, and in particular another state or other authority which has the legal powers to stop or prevent the infringement does not, or cannot, use them to that effect;
- d) the action taken by way of self-defence is limited to what is necessary to stop or prevent the infringement, i.e., to the needs of defence.

When are these conditions met? The facts and circumstances in any given situation may not lead to clear determinations—indeed, the threatened party is likely to have a rather different perception of such facts and circumstances than the threatening state.

The mere fact that Zendia possesses destructive capabilities that could be used against Ruritania cannot be sufficient to indicate imminent attack—

⁵ *Oppenheim’s International Law: Ninth Edition*, 1991, p. 412.

BOX 7.1 Self-defense and Self-help

Article 51 acknowledges the right of a nation to engage in the use of armed force for self-defense, including the situation in which the nation is the target of an armed attack, even without Security Council authorization. (The issue of whether a nation may respond militarily without Security Council authorization if it is the target of a use of force short of an armed attack is less clear, with evidence to support both sides of this position.¹) Although the term “self-defense” is undefined in the UN Charter, it is convenient to consider three different types of actions, all of which involve the use of force in response to an attack.

- A Type 1 action is a use of force taken to halt or curb an attack in progress or to mitigate its effects. Type 1 actions do not apply after the attack ceases, because all of the harm that the attack can cause has already been caused at that point.
- A Type 2 action is a use of force in which a nation is the first to use force because it has good reason to conclude that it is about to be attacked and that there is no other alternative that will forestall such an action. Type 2 actions are sometimes called actions of anticipatory self-defense.²
- A Type 3 action is a use of force aimed at reducing the likelihood that the original attacker will continue its attacks in the future. Type 3 actions are predicated on the assumption that the original attacker has in mind a set of attacks, only one of which has occurred, and can be regarded as a kind of anticipatory self-defense against these likely future attacks. An example of a Type 3 action is the 1986 El Dorado Canyon bombing on Libya, which was justified as an act of self-defense against a continuing Libyan-sponsored terrorist threat against U.S. citizens.³ (Note that under domestic law as it applies to private persons, Type 3 actions are generally not legal, though Type 1 actions taken in self-defense are sometimes justified under common law, as indicated in Section 5.2.)

Many nations, including the United States, have asserted rights under the UN Charter to all three types of action under the rubric of self-defense. At the same time, other nations (especially including the target of such action) have claimed that a Type 3 action is really an illegal reprisal—that is, an act of punishment or revenge.

In the context of cyberattack and active defense, a Type 1 action corresponds to active threat neutralization—a cyberattack launched in response to an incoming cyberattack that is intended to neutralize the threat and to stop further damage from occurring. A Type 3 action corresponds to a cyberattack that is intended to dissuade the attacker from launching further attacks in the future.

The difference between Type 1 and Type 3 actions is significant because a Type 3 action is technically easier to conduct than a Type 1 action under some circumstances. For example, it may easily come to pass that an incoming cyberattack can be identified as emanating from Zendia and that the Zendian national

authorities should be held responsible for it. A Type 3 action could then take the form of any kind of attack, cyber or kinetic, against Zendia—without the enormous difficulty of identifying a specific access path to the controllers behind the attack (necessary for a Type 1 action). In addition and depending on the circumstances, a Type 1 action could be followed by a Type 3 action. That is, a policy decision might be made to take a Type 3 action to ensure that no more hostile actions were taken in the future.

Self-defense actions are clearly permissible when a nation or its forces have experienced an armed attack. Under standing rules of engagement, a missile fired on a U.S. fighter plane or a fire-control radar locked on the airplane would count as an armed attack, and self-defense actions (e.g., bombing the missile site or the radar) would be allowable. In a similar vein, cyberattacks that compromise the ability of units of the DOD to perform the DOD's mission might well be regarded as an armed attack, and indeed STRATCOM has the authority to conduct response actions to neutralize such threats (Chapter 3).

If a nation has been the target of a use of force (a cyberattack) that does not rise to the threshold of an armed attack, responses made by the victimized nation fall into the category of self-help. What self-help actions are permissible under the UN Charter?

Certainly any action that does not amount to a use of force is legal under the UN Charter as long as it does not violate some existing treaty obligation. An example of such an action might well be non-cooperative but non-destructive intelligence gathering about the attacking system. In addition, a small-scale Type 1 action to neutralize an incoming cyberattack aimed at a single system is likely to be permissible. (An analogy from physical space might be the small-scale use of force to shoot armed border crossers.)

¹ Department of Defense, Office of General Counsel, *An Assessment of International Legal Issues in Information Operations, Second Edition*, November 1999.

² See, for example, *Oppenheim's International Law: Ninth Edition*, 1991, p. 412.

³ The raid was the culmination of increasing tensions between the United States and Libya. Since 1973, Muammar Qadhafi asserted Libyan control over the Gulf of Sidra, a claim not recognized under international law (which recognizes only a 12-mile-from-shore claim for national waters). In 1981, the United States conducted naval exercises in the area claimed by Libya, with the result that two Libyan fighter-bombers sent to challenge the United States presence were shot down. Tensions continued to increase, and in March 1986, Libya launched six SA-5 missiles against the U.S. Sixth Fleet, then operating nearby in the Mediterranean. In subsequent action, the United States destroyed two Libyan vessels. In early April 1986, a bomb exploded in a Berlin discotheque, killing a U.S. soldier and injuring 63 U.S. soldiers, among others. The United States asserted that it had communications intercepts proving Libyan sponsorship of the bombing, and Operation El Dorado Canyon occurred shortly thereafter, as the United States had at the time no reason to expect such attacks to cease. In May 2001, Qadhafi acknowledged to a German newspaper that Libya had been behind the discotheque bombing 15 years earlier, which was carried out apparently in retaliation for the U.S. sinking of the two vessels in March 1986.

otherwise, the mere existence of armed forces of an adversary would be sufficient justification. But if Zendia can use these capabilities effectively against Ruritania and with serious consequences without warning, and Zendia has indicated hostile intent toward Ruritania in other (perhaps non-military) ways, outside observers may indeed be more likely to judge that the conditions for anticipatory self-defense have been met.

7.2.1.2 *Jus in Bello*

Once armed conflict has begun, the conduct of a nation's armed forces is subject to a variety of constraints. *Jus in bello* is governed largely by the Hague Conferences of 1899 and 1907, the Geneva Conventions, and customary international law.

- *Military necessity.* Valid targets are limited to those that make a direct contribution to the enemy's war effort, or those whose damage or destruction would produce a military advantage because of their nature, location, purpose, or use. Thus, enemy military forces (and their equipment and stores) may be attacked at will, as is also true for civilians and civilian property that make a direct contribution to the war effort. Assets that do not contribute to the war effort or whose destruction would provide no significant military advantage may not be deliberately targeted by cyber or kinetic means. LOAC also provides for a category of specially and (in theory) universally protected facilities such as hospitals and religious facilities.

- *Proportionality.* It is understood that attacks on valid military targets may result in collateral injury and damage to civilian assets or people. Some degree of collateral damage is allowable, but not if the foreseeable collateral damage is disproportionate compared to the military advantage likely to be gained from the attack. In the event that military and nonmilitary assets are circumstantially commingled (e.g., the use of a common electric grid to power both military and civilian facilities), the attacker must make a proportionality judgment. But in instances when the enemy has deliberately intermingled military and non-military assets or people, the enemy must then assume some responsibility for the collateral damage that may result.

Put differently, LOAC always obligates a would-be attacker to make reasonable proportionality judgments. What is less clear, and may depend on circumstances, are the conditions under which the enemy has a legal responsibility to refrain from deliberately commingling military assets with non-military assets or more generally to separate such assets. For example, the enemy may have deliberately placed "human shields" around military targets. In such a case, the enemy is clearly in violation

of LOAC and bears the responsibility for any injury to the hostages if the target is attacked. However, in an extreme case where the likely deaths and injuries among the hostages are disproportionate to the military advantage to the attacker, the attacker is obligated to take into account the presence and likely deaths of those human shields in making a proportionality judgment about a possible attack.

A common misperception about proportionality as a rule of *jus in bello* is that it requires the victim of an attack to respond only in ways that cause the original attacker approximately the same amount or degree of pain that the victim experienced. This kind of response is generally characterized as a commensurate response, and although commensuration and commensurate response are often used by policy makers as guideposts in formulating responses to external attack, they are not required by LOAC.

- *Perfidy*. Acts of perfidy seek to deceive an enemy into believing that he is obligated under the law of armed conflict to extend special protection to a friendly asset when such is not the case. For example, by convention and customary law, certain persons and property may not be legitimately attacked, including prisoners of war and prisoners-of-war camps, the wounded and sick, and medical personnel, vehicles, aircraft, and vessels. Persons and property in this category must be identified with visual and electronic symbols, and misuse of these symbols to prevent a legitimate military target from being attacked constitutes the war crime of perfidy. In addition, it is unlawful to feign surrender, illness, or death to gain an advantage in combat, or to broadcast a false report that both sides had agreed to a cease-fire or armistice. At the same time, ruses of war are explicitly permissible. A ruse of war is intended to mislead an adversary or to induce him to act recklessly but its use infringes no rule of international law applicable in armed conflict and does not mislead the adversary into believing that he is entitled to special protection. Camouflage, decoys, mock operations, and misinformation are all permitted ruses.

- *Distinction*. Distinction requires armed forces to make reasonable efforts to distinguish between military and civilian assets and between military personnel and civilians, and to refrain from deliberately attacking civilians or civilian assets. However, there are two important classes of civilians or civilian assets—those that have been compromised and used (illegally) to shield the actions of a party to the conflict and those that suffer inadvertent or accidental consequences (“collateral damage”) of an attack. Responsibility for harm is apportioned differently depending on the class to which a given civilian or civilian asset belongs (Box 7.2).

- *Neutrality*. A nation may declare itself to be neutral, and is entitled to immunity from attack by either side at war, as long as the neutral nation does not assist either side militarily and acts to prevent its territory

BOX 7.2 Avoiding Harm to Innocent Parties

The principle of distinction requires military forces to minimize harm to innocent parties—that is, non-combatants that are not actively engaged in helping to prosecute the war. But three categories of “innocent parties” must be distinguished, especially in the cyber context.

- *Category A*—An innocent party that is compromised by an adversary and then used to shield the adversary’s actions. For example, an adversary (Zendia) that uses human civilians as shields to protect its anti-aircraft sites is using this kind of innocent party. Zendia would also be doing so if it launched a cyberattack against Ruritania through the use of a compromised and innocent third-party computer (e.g., one belonging to civilians).
- *Category B*—An innocent party that is caught up in some effect that was unpredicted or could not have been expected. For example, a Zendian civilian truck in the desert is struck inadvertently by the empty drop tanks of a Ruritanian fighter-bomber en route to its target, and all those inside the truck are killed. Or, a Ruritanian cyberattack strikes a Zendian generator powering the Zendian ministry of defense, leading to a cascading power failure that disables hospitals in which Zendian patients then die.
- *Category C*—An innocent party that is granted special protection under the Geneva Convention, such as a hospital, and is then used as a facility from which to launch attacks. For example, the Zendian adversary that places mortars on the roof of a hospital is using Category C innocent parties. Or, Zendia launches a cyberattack on Ruritania using the servers and Internet connections of a Zendian hospital.

Distinguishing between these kinds of innocent parties is important because the categories of parties harmed have different implications for responsibility. If

from being so used. Accordingly, there exists a right for a threatened state “to use force to neutralize a continuing threat located in the territory of a neutral state, but not acting on its behalf, when the neutral state is unable or unwilling to fulfill its responsibility to prevent the use of its territory as a base or sanctuary for attacks on another nation.”⁶ Note also that under item 3 of UN Security Council Resolution 1368 (adopted on September 12, 2001),⁷ which calls on all member states “to work together urgently to bring to justice the perpetrators, organizers and sponsors of these terrorist attacks” and stresses that “those responsible for aiding, supporting, or harboring the perpetrators, organizers and sponsors of these acts will be

⁶ Department of Defense, Office of General Counsel, *An Assessment of International Legal Issues in Information Operations*, Second Edition, November 1999.

⁷ United Nations Security Council Resolution 1368 (2001), accessed at <http://daccessdds.un.org/doc/UNDOC/GEN/N01/533/82/PDF/N0153382.pdf?OpenElement>.

Category A innocent parties are harmed, some responsibility attaches to Zendia for placing innocent parties in harm's way. Some degree of responsibility may attach to Ruritania if the attack did not meet the requirements of proportionality—that is, if the military value of the target shielded was small by comparison to the loss of Zendian civilian life. If Category B innocent parties are harmed, the responsibility does not fall on Zendia, and if Ruritania took reasonable care in route planning, no responsibility attaches to Ruritania either. If Category C innocent parties are harmed, the legal responsibility for those consequences falls entirely on the Zendian adversary under LOAC.

In active defense scenarios calling for threat neutralization, there are many valid concerns about a counterstrike that does harm to some innocent party. But at least in some scenarios involving innocent third-party computers (that is, in Category A), a Ruritanian response against those compromised computers could be conducted within the bounds of LOAC, and the harm resulting to those third parties would be the responsibility of Zendia and not Ruritania.

Of course, Ruritania would have to address several other concerns before feeling confident in the legality and wisdom of a counterstrike. First, even if a counterstrike is entirely legal, it may come with other costs, such as those associated with public opinion or ethical considerations. If a counterstrike disables the hospital computer and deaths result, there may be censure for Ruritania, even if the counterstrike was within Ruritania's legal rights to conduct. Second, Ruritania would have to take reasonable care to determine that the incoming cyberattack was indeed coming from the computer in question, because Zendia might have also planted evidence so as to prompt a counterstrike against a computer that was not involved in the attack at all. Third, Ruritania would still have to make reasonable efforts to ensure that its attack on the hospital computer did not have unintended cascading effects (e.g., beyond the particular node on the hospital network from which the attack was emanating).

held accountable," and under related developments in international law, even neutral states have affirmative obligations to refrain from harboring perpetrators of terrorist attacks. The United States has asserted the right of self-defense in this context on a number of occasions, including the 1998 cruise missile attack against a terrorist training camp in Afghanistan and a chemical plant in Sudan in which the United States asserted that chemical weapons had been manufactured; the 1993 cruise missile attack against the Iraqi intelligence service headquarters which the United States held responsible for a conspiracy to assassinate President George H.W. Bush; and the 1986 bombing raid against Libya in response to Libya's continuing support for terrorism against U.S. military forces and other U.S. interests.

- *Discrimination.* Nations have agreed to refrain from using certain weapons, such as biological and chemical weapons, at least in part

because they are inherently indiscriminate weapons (that is, they cannot be directed against combatants only). However, there is no ban as such on indiscriminate weapons per se—the harm to non-combatants is minimized through adherence to requirements of proportionality.

It is worth emphasizing that *jus ad bellum* and *jus in bello* are two different bodies of law, applicable at different times. Once armed conflict has started (whether or not *jus ad bellum* has been followed in the starting of that conflict), *jus in bello* is the body of law that applies.

7.2.2 Applying the Law of Armed Conflict to Cyberattack

This section addresses some of the issues that might arise in applying international law to cyberattack. Some issues arise when a nation is the target of a cyberattack and must consider legal issues in formulating an appropriate and effective response—and its decision depends on (among other things) whether it is in an ongoing state of hostilities with the perpetrator of that cyberattack. Other issues arise when a nation may wish to launch a cyberattack against another party prior to the outbreak of hostilities but without intending to give the other side a legal basis for regarding its action as starting a general state of hostilities.⁸ Still other issues arise when cyberattack is conducted in the context of an ongoing conflict—that is, while hostilities are in progress. And a different set of standards and legal regimes may govern responses to cyberattacks launched by non-state actors.

To be fair, many or most of the same issues addressed below arise when kinetic weapons are used in conflict. But cyberweapons are newer and have certain characteristics not shared with kinetic weapons, which implies that fewer precedents and analyses are available and that the application of LOAC principles may not be as straightforward as they are when kinetic weapons are involved.

On the broad question regarding cyberattack, the committee starts with two basic premises that guide subsequent discussion:⁹

⁸ Department of Defense, Office of General Counsel, *An Assessment of International Legal Issues in Information Operations*, Second Edition, November 1999.

⁹ These points are addressed in a number of legal analyses, including Michael Schmitt, "Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework," *Columbia Journal of Transnational Law* 37:885-937, 1999; Duncan B. Hollis, "New Tools, New Rules: International Law and Information Operations," pp. 59-72 in *Ideas As Weapons: Influence and Perception in Modern Warfare*, G. David and T. McKeldin, eds., Potomac Books, Inc., 2009; and Jason Barkham, "Information Warfare and International Law on the Use of Force," *New York University International Law and Politics* 34:57-113, 2001. Schmitt's and Hollis's analyses are summarized in Appendix D.

- Cyberattack cannot be regarded as a more “benign” form of warfare or as always falling short of “armed attack” or “use of force” simply because a cyberattack targets computers and networks. The magnitude, scale, and nature of a cyberattack’s effects, both direct and indirect, have to be taken into account in ascertaining its significance, and it is not simply the modality of the attack that matters.¹⁰

- Despite the fact that cyberattack is a relatively new form of weapon, acknowledged armed conflict involving the use of cyberweapons is subject to LOAC and UN Charter law. That is, LOAC’s precepts regarding *jus ad bellum* and *jus in bello* continue to have validity in a cyberattack context. Nevertheless, because of the novelty of such weapons, there will be uncertainties in how LOAC and UN Charter law might apply in any given instance. An effects-based analysis suggests that the ambiguities are fewest when cyberattacks cause physical damage to property and loss of life in ways that are comparable to kinetic attacks and traditional war is involved, because traditional LOAC provides various relevant precedents and analogies. The ambiguities multiply in number and complexity when the effects do not entail physical damage or loss of life but do have other negative effects on another nation.¹¹

Appendix D summarizes several other views on cyberattack as a use of force.

Also, as Hollis notes,¹² traditional LOAC and the UN Charter are largely silent on how to address conflict involving non-state actors, even though non-state actors (in particular, terrorist groups) are playing larger roles in the security environment today. This point is addressed in Section 7.2.3.1 (on terrorists), Section 7.2.3.2 (on multinational corporations), and Section 7.2.3.3 (on individuals).

7.2.2.1 Prior to the Outbreak of Hostilities—Applying *Jus ad Bellum*

An important question of *jus ad bellum* in this report is whether, or more precisely, when, a given cyberattack constitutes a “use of force” or an “armed attack.” But as a number of analysts have noted,¹³ the relevant

¹⁰Department of Defense, Office of General Counsel, *An Assessment of International Legal Issues in Information Operations*, Second Edition, November, 1999.

¹¹See Jason Barkham, “Information Warfare and International Law on the Use of Force,” *New York University International Law and Politics* 34:57-113, 2001.

¹²Duncan B. Hollis, “New Tools, New Rules: International Law and Information Operations,” pp. 59-72 in *Ideas As Weapons: Influence and Perception in Modern Warfare*, G. David and T. McKeldin, eds., Potomac Books, Inc., 2009.

¹³Michael Schmitt, “Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework,” *Columbia Journal of Transnational Law* 37:885-

question is not so much whether a cyberattack constitutes a “use of force” but rather whether a cyberattack *with a specified effect* constitutes a “use of force.” That is, the effects of a given cyberattack are the appropriate point of departure for an analysis of this question, rather than the specific mechanism used to achieve these effects.

7.2.2.1.1 *The Uncertainties in Identification and Attribution*

Application of LOAC in a cyber context requires identification of the party responsible for an act of cyber aggression. But as noted in Chapter 2, it may be difficult even to know when a cyberattack has begun, who the attacker is, and what the purpose and effects of the cyberattack are/were. Indeed, it may be difficult to identify even the nature of the involved party (e.g., a government, a terrorist group, an individual), let alone the name of the country or the terrorist group or the individual. Knowing the nature of the party is an important element in determining the appropriate response.¹⁴ And, of course, knowing which country, terrorist group, or individual is in fact responsible is essential if any specific response involving attack is deemed appropriate.

- What, if any, is the responsibility of an attacking nation to ascertain the physical location of a computer or network that it attacks? Where kinetic weapons are involved, attacking a particular target requires knowledge of the target’s physical location. But it is often possible for a cyberweapon to attack a target whose location is known only as an IP address or some other machine-readable address that does not necessarily correspond to a specific or a known physical location. Yet physical location may matter (a point that relates to notions of territorial integrity) in determining whether a given cyber target belongs to or is under the control of an adversary.

- What degree of certainty about the identity of an attacker is needed legally before a cyberattack may be launched to neutralize it? How, if at all, does this differ from what is needed for policy purposes?

Box 7.3 provides some scenarios in which such questions arise.

937, 1999; Jason Barkham, “Information Warfare and International Law on the Use of Force,” *New York University Journal of International Law and Politics* 34:57-113, 2001; Department of Defense, Office of General Counsel, *An Assessment of International Legal Issues in Information Operations*, Second Edition, November 1999. An exposition by Brownlie in 1963 discusses a “results-oriented” approach, but of course without reference to cyberattack per se. See Ian Brownlie, *International Law and the Use of Force by States*, 1963.

¹⁴ Sections 2.4.2 and 2.4.3 describe some of the issues involved in making such a determination.

BOX 7.3 Uncertainties in Identification and Attribution—Possible Examples

The following examples illustrate possible scenarios in which uncertainties in identification and attribution arise.

- *During conflict between the United States and Zendia, a U.S. cyberattack is launched on a computer controlling a Zendian air defense network. A normally reliable human informant passes on a message to the United States, but the message is unfortunately incomplete, and the only information passed along is the computer's electronic identifier, such as an IP address or a MAC (Media Access Control) address; its physical location is unknown. The open question is whether this computer is a valid military target for a U.S. cyberattack and the extent to which the United States has an obligation to ascertain its physical location prior to such an attack.*
- *During a time of international tension (say, U.S. forces are on an elevated alert status), the United States experiences a cyberattack on its military communications that is seriously disruptive. The United States must restore its communications quickly but lacks the intelligence information to make a definitive assessment of the ultimate source of the attack. The open question is whether it can lawfully act against the proximate sources of the attack in order to terminate the threat and restore its communications capability, even though it is by no means certain that the "proximate source" is actually the ultimate source and may simply have been exploited by the ultimate source. (A proximate source might be a neutral nation, or a nation whose relations with the United States are not particularly good. If the latter, a U.S. attempt to neutralize the attack might thus exacerbate tensions with that nation.)*

One practical consequence of these uncertainties is that a nation seeking UN action in response to a cyberattack would be unlikely to see rapid action, since much of the necessary information might not be available promptly. (Indeed, consider as a benchmark the history of long and extended Security Council debate on authorizations for armed conflict involving kinetic force.)

7.2.2.1.2 Criteria for Defining "Use of Force" and "Armed Attack"¹⁵

Traditional LOAC emphasizes death or physical injury to people and destruction of physical property as criteria for the definitions of "use of force" and "armed attack." But modern society depends on the existence

¹⁵ A related perspective can be found in Jason Barkham, "Information Warfare and International Law on the Use of Force," *New York University International Law and Politics* 34:57-113, 2001.

and proper functioning of an extensive infrastructure that itself is increasingly controlled by information technology. Actions that significantly interfere with the functionality of that infrastructure can reasonably be regarded as uses of force, whether or not they cause immediate physical damage. Thus, cyberattacks on the controlling information technology for a nation's infrastructure that had a significant impact on the functioning of that infrastructure (whether or not it caused immediate large-scale death or destruction of property) would be an armed attack for Article 51 purposes, just as would a kinetic attack that somehow managed to shut down the system without such immediate secondary effects.

How far would a cyberattack on a nation's infrastructure have to go before it was regarded as a use of force or an armed attack? Scale of effect is one important factor in distinguishing between an armed attack and a use of force. For example, an armed attack would presumably involve a use of force that resulted in a large scale of effect. It is unclear if there are other differentiating factors in addition to scale of effect. (Neither "armed attack" nor "use of force" necessarily requires the use of traditional kinetic weapons.)

Schmitt's examples of cyberattacks that do and do not qualify as a use of force are useful for establishing a continuum of scale.¹⁶ At one end, Schmitt argues that a cyberattack on an air traffic control system resulting in a plane crash and many deaths clearly does qualify as a use of force, whereas a computer network attack on a single university computer network designed to disrupt military-related research occurring in campus laboratories does not. In between these two ends of the spectrum are a number of problematic cases (Box 7.4) that raise a number of questions.

- What is the minimum length of time, if any, for which a serious disruption to critical infrastructure must last for it to be regarded as a use of force or an armed attack? (This is not to say that time is the only variable involved in such an assessment.)
- Under what circumstances, if any, can a non-lethal and continuing but reversible cyberattack that interferes with the functionality of a target network (e.g., against a photo reconnaissance satellite) be regarded as a use of force or an armed attack?
- Under what circumstances, if any, can a cyberattack (e.g., against a stock market's data, against a factory process) whose disruptive but not actually destructive effects build slowly and gradually be regarded as a use of force or an armed attack?

¹⁶ Michael Schmitt, "Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework," *Columbia Journal of Transnational Law* 37:885-937, 1999.

BOX 7.4 Cyberattacks as a Possible “Use of Force”

The following examples illustrate possible scenarios that raise questions about the appropriate definition of the “use of force.”

- *A cyberattack temporarily disrupts Zendia’s stock exchanges and makes trading impossible for a short period.* Bombs dropped on Zendia’s stock exchanges (at night, so that casualties were minimized) would be regarded as a use of force or an armed attack by most observers, even if physical backup facilities were promptly available so that actual trading was disrupted only for a short time (e.g., a few hours). The posited cyberattack could have the same economic effects, except that the buildings themselves would not be destroyed. In this case, the cyberattack may be less likely to be regarded as a use of force than a kinetic attack with the same (temporary) economic effect, simply because the lack of physical destruction would reduce the scale of the damage caused. However, a cyberattack against the stock exchanges that occurs repeatedly and continuously, so that trading is disrupted for an extended period of time (e.g., days or weeks), would surely constitute a use of force or even an armed attack, even if no buildings were destroyed.
- *A cyberattack is launched against the ground station of a Zendian military photo-reconnaissance satellite.* Neither the satellite nor the ground station is physically damaged, but Zendia is temporarily unable to download imagery. The open question is whether such an act might plausibly be interpreted as a use of force, based on the argument that the inability to download imagery might be a prelude to an attack on Zendia, even if no (permanent) damage has been done to Zendia.
- *A cyberattack has effects that build slowly and gradually.* For example, a cyberattack against a stock exchange might corrupt the data used to make trades. Again, no physical damage occurs to buildings, and in addition trading continues, albeit in a misinformed manner. Over time, the effects of such an attack could wreak havoc with the market if continued over that time.¹ If and when the effects were discovered, public confidence in the market could well plummet, and economic chaos could result. An open question is the degree of economic loss, chaos, and reduction in public confidence that would make such an attack a use of force.
- *A cyberattack is aimed at corrupting a manufacturing process.* In this scenario, the manufacturing process is altered in such a way that certain flaws are introduced into a product that do not show up on initial acceptance testing but manifest themselves many months later in the form of reduced reliability, occasional catastrophic failure, significant insurance losses, and a few deaths. Here, one open question relates to the significance of the effects of the attack, recognizing the “boiling the frog” phenomenon—a sudden change may be recognized as significant, but a gradual change of the same magnitude may not be.

¹ As a demonstration that slowly accumulating error can have large consequences, consider that the Vancouver stock exchange index introduced in 1982 was undervalued by 48 percent 22 months later compared to its “true” value—the reported value of the index was 524.881, whereas the correctly calculated value should have been 1009.811. This discrepancy was the result of roundoff error, accumulated over time. See B.D. McCullough and H.D. Vinod, “The Numerical Reliability of Econometric Software,” *Journal of Economic Literature* 37(2):633-665, June 1999.

The question of scale above points to a more general problem—the inability to distinguish at the point of discovery that a cyberattack is taking place between one that seeks to cause large-scale damage (which would almost certainly constitute an armed attack) and one that seeks to cause only very limited damage (which might constitute a use of force if not an armed attack). The problem of a nation figuring out when a given act that may appear to be hostile is—or is not—a precursor to more serious hostile actions that will create additional damage is not unique to cyberattack, as illustrated by the Tonkin Gulf incident (in which the United States was arguably too quick to see a grave provocation) and Stalin’s refusal to believe reports of Nazi preparations and initial incursions in June 1941. Similarly, an airplane penetrating a nation’s airspace without authorization may simply be off course, or it may instead be carrying nuclear weapons with hostile intent. The nation in question has an obligation to try to determine if the airplane represents a true threat, but it surely has a right to shoot down the airplane if it reasonably makes such a determination. The open question is what the nation can do if it is uncertain about the threatening nature of the airplane.

Although waiting to see what the attack does is the only certain way to determine the scale and extent of its effects, waiting may not be a viable option for decision makers when they are notified that a cyberattack on their nation is underway. In addition, leaders of a state often wish to calibrate a response to an attack to be of the same scale as that attack—and if decision makers do not know the scale of the attack, how are they to calibrate a response?¹⁷

The scale question also raises the issue of whether there is, or should be, a class of “hostile” cyber actions (that is, certain kinds of cyberattack) that are recognized as not so immediately destructive as to be clear acts of “uses of force” or “armed attack,” but that nonetheless entitle the target to some measure of immediate real-time response—commensurate self-defense—that goes beyond just trying to protect the immediate target. (Such a regime might have some counter-escalation effects, because a potential aggressor would not be assured of immunity from a response from its victim.) A regime designed with an overriding priority to discourage escalation of cyberconflict would not recognize the existence of such a class but rather obligate the target to accept the initial consequences of those hostile cyber actions and respond (whether by force or otherwise) only afterward.

¹⁷ Jason Barkham, “Information Warfare and International Law on the Use of Force,” *New York University International Law and Politics* 34:57-113, 2001.

7.2.2.1.3 Definition of "Threat of Force"

Article 2(4) prohibits nations from threatening the use of force. When the coercive instruments are traditional weapons, a threat generally takes the form of "We will do destructive act X if you do not take action Y (that is, trying to compel the adversary to take action Y) or if you do take action Z (that is, trying to deter the adversary from taking action Z)."

- Does a threat to use existing vulnerabilities in an adversary computer system or network constitute a threat of the use of force under the UN Charter? Because an existing vulnerability can be used for cyberattack (which can be a use of force) or cyberexploitation (which is not considered a use of force, as discussed in Section 7.2.2.1.5), the answer is not clear.
- Does it matter how those vulnerabilities got there? Does introducing vulnerabilities into an adversary's system or network constitute a threat of force, especially if they remain unused for the moment?

Box 7.5 provides examples illustrating how such questions might arise.

7.2.2.1.4 Distinctions Between Economic Sanctions and Blockades¹⁸

Under international law, economic sanctions appear not to constitute a use of force, even if they result in death and destruction on a scale that would have constituted a use of force if they were caused by traditional military forces, although this interpretation is often questioned by the nation targeted by the sanctions. Article 41 of the UN Charter gives the Security Council authority to decide what measures count as "not involving the use of armed force," and it explicitly recognizes that measures not involving the use of armed force include the "complete or partial interruption of economic relations."¹⁹

In this instance, international law does appear to differentiate between different means used to accomplish the same end. That is, economic sanctions and blockades could easily result in similar outcomes, but there are two key differences. First, sanctions are, by definition, a refusal of participating nations to trade with the targeted party, either unilaterally (by virtue of a national choice) or collectively (by virtue of agreement to adhere to UN mandates regarding sanctions). That is, sanctions involve refraining from engaging in a trading relationship that is not obligatory. By contrast, blockades interfere with trade involving any and all parties,

¹⁸ See also Jason Barkham, "Information Warfare and International Law on the Use of Force," *New York University International Law and Politics* 34:57-113, 2001, whose analysis roughly parallels the argument of this subsection.

¹⁹ See <http://www.un.org/aboutun/charter/chapter7.htm>.

BOX 7.5 Threats of Force—Possible Examples

The following examples illustrate possible scenarios that raise questions about the definition of the “threat of force.”

- *Zendia introduces cyber vulnerabilities into the critical infrastructure of its adversary Ruritania, but does not take advantage of them.* Since Ruritania suffers no ill effects from the fact that its infrastructure now has a number of vulnerabilities, no armed attack or even use of force has occurred. Ruritania learns of the Zendian penetration because its cybersecurity experts have detected it technically. Does the Zendian action of introducing cyber vulnerabilities constitute a threat of force against Ruritania? Does it make a difference if these vulnerabilities could be used equally well for cyberexploitation as for cyberattack? Does the possibility that Zendia could take advantage of those agents on a moment’s notice make a cyberattack on Ruritania imminent, and if so, does it justify a Ruritanian strike on Zendia (cyber or otherwise) as an act of anticipatory self-defense?

A possibly helpful analogy is that of digging a tunnel underneath a border that terminates underneath a military facility. If Zendia digs such a tunnel under the Zendia-Ruritania border, and Ruritania discovers it, Ruritania may well regard it as a hostile act. But whether the tunnel amounts to an indication of imminent hostilities that would justify a Ruritanian strike on Zendia depends on many other factors.

- *Zendia discovers cyber vulnerabilities in the critical infrastructure of its adversary Ruritania, but does not take advantage of them.* These vulnerabilities are found in software used by both Zendia and Ruritania and supplied by a third-nation vendor. If Zendia notifies Ruritania of these vulnerabilities during a time of tension between the two nations, has Zendia threatened to use force against Ruritania?

willing and unwilling. Second, effective economic sanctions generally require coordinated multilateral actions, whereas blockades can be conducted unilaterally, though the coordination mechanism may or may not be tied to UN actions.²⁰

From the standpoint of effects-based analysis, traditional LOAC thus has some inconsistencies embedded within it regarding means used for

²⁰ Some economic sanctions can be imposed unilaterally and still be effective. For example, if the Zendian armed forces use a sophisticated weapons system that was originally produced in the United States, spare parts for that system may only be available from the United States. The United States could unilaterally choose to refrain from selling spare parts for that system to Zendia without violating LOAC, and such an action could have significant effects on the Zendian armed forces as the weapons system deteriorated due to a lack of spare parts. In addition, multilateral sanctions need not necessarily involve the United Nations, as demonstrated by the Arab boycott of Israel, the Arab oil embargo of 1973, and the 2008 financial sanctions against Iran.

economic coercion, even if cyberattack is not involved. At the very least, it draws distinctions that are not entirely clear-cut. Accordingly, it is not surprising that such inconsistencies might emerge if cyberattack is used for economic coercion without the immediate loss of life or property. Legal analysts must thus determine the appropriate analogy that should guide national thinking about cyberattacks that result in severe economic dislocation. In particular, are such cyberattacks more like economic sanctions or a blockade (or even some form of kinetic attack, such as the mining of a harbor)?

This question is particularly salient in the context of Internet-enabled commerce. The UN Security Council could decide to impose economic sanctions on a nation in order to compel that nation to follow some directive, and in principle those sanctions can be quite broad and sweeping. If a large part of the target nation's commerce was enabled through international Internet connections, the omission of such commerce from the sanctions regime might be a serious loophole.²¹ On the other hand, cyberattacks against the target nation might be required to prevent such commerce from taking place in a manner analogous to the UN's use of naval and air forces to enforce certain past economic sanctions.

Box 7.6 provides some scenarios in which the question of the most appropriate analogy arises.

One last caveat regarding the economic dimension of cyberattack: It is possible to imagine cyberattack as a tool for pursuing goals related to economic competition and/or economic warfare. It is clear that the laws of armed conflict and the UN Charter prohibit the use of force—cyber as well as kinetic force—in pursuit of purely economic or territorial gain. But the legitimacy of cyberattacks that do not constitute a use of force for economic gain is not entirely clear. (As noted in Section 2.6.2, some nations do conduct espionage for economic purposes (an activity not prohibited by international law), and cyberattack might well be used to conduct espionage. And, as noted in Section 4.2.2, destructive cyberattacks might be used to gain economic advantage.)

7.2.2.1.5 *The De Facto Exception for Espionage*

Espionage is an illegal activity under the domestic laws of virtually all nations, but not under international law. For example, Hays Parks has written:

²¹ As a practical matter, many of the nations that are subject to sanctions are often not heavily dependent on Internet commerce, or at least they are not today. In addition, sanctions are often not generalized but rather are targeted at specific goods such as arms.

BOX 7.6 Cyberattack as Blockades or Sanctions— Possible Examples

The following examples illustrate possible scenarios that raise questions about whether to treat a cyberattack as a blockade or an economic sanction.

- *A continuing cyberattack that effectively disconnects Zendia's access to the global Internet, when Zendia is the target of UN economic sanctions.* In the modern era, a nation's economic relations with the outside world may be more dependent on the Internet than a nation was dependent on maritime shipping in the mid-20th century. Should this type of cyberattack—perhaps performed openly by a permanent member of the UN Security Council—be regarded as a blockade enforced through electronic means or as the enforcement of economic sanctions? Does it matter if the cyberattack targets only the Zendian connections to the outside world versus targeting internal communications nodes and routers?
- *A cyberattack that shuts down a key industry or segment of the armed forces of the targeted nation.* Economic sanctions and blockades can be narrowly tailored to affect only certain industries. For example, sanctions and blockades could prevent the sale or distribution of spare parts necessary for the continuing operation of a certain industry. The same is true for spare parts needed to maintain and operate certain weapons systems. But a cyberattack could also have similar effects—and in particular could be carried out in such a way that the industry or military segment targeted was degraded slowly over time in a manner similar to its degradation due to the lack of spare parts. Thus, this kind of cyberattack could have effects identical to that of either blockades or economic sanctions, though one is regarded as a use of force and the other not.

Each nation endeavors to deny intelligence gathering within its territory through domestic laws Prosecution under domestic law (or the threat thereof) constitutes a form of denial of information rather than the assertion of a *per se* violation of international law; domestic laws are promulgated in such a way to deny foreign intelligence collection efforts within a nation's territory without inhibiting that nation's efforts to collect intelligence about other nations. No serious proposal has ever been made within the international community to prohibit intelligence collection as a violation of international law because of the tacit acknowledgement by nations that it is important to all, and practiced by each.²²

²² W. Hays Parks, "The International Law of Intelligence Collection," pp. 433-434 in *National Security Law*, John Norton Moore et al., eds., 1990, cited in Roger D. Scott, "Territorially Intrusive Intelligence Collection and International Law, *Air Force Law Review* 46:217-226, 1999, available at [http://permanent.access.gpo.gov/lps28111/Vol.46 \(1999\)/scottfx4\[1\].doc](http://permanent.access.gpo.gov/lps28111/Vol.46%20(1999)/scottfx4[1].doc).

If this legal approach is accepted, espionage conducted by or through the use of a computer—that is, cyberexploitations—is permissible under the LOAC regime, even if techniques are used that could also be used for destructive cyberattack. For example, cyberattacks may be used to disable cybersecurity mechanisms on a computer of interest so that a keystroke monitor can be placed on that computer.

Nevertheless, espionage may raise LOAC issues if a clear distinction cannot be drawn between a given act of espionage and the use of force. For example, Roger Scott notes that certain forms of espionage—for instance involving ships, submarines, or aircraft as the collection platforms—have indeed been seen as military threats and have been treated as matters of armed aggression permitting a military response rather than domestic crimes demanding a law enforcement response.²³

One common thread here appears to be that the collection platform is or appears to be a military asset—a plane, a ship, a submarine—that could, in principle, conduct kinetic actions against the targeted nation. In all of these cases, the question of intent is central to the targeted nation at the time the potentially hostile platform is detected. Furthermore, the distinction between a cyberattack and a cyberexploitation may be very hard to draw from a technical standpoint, since both start with taking advantage of a vulnerability.

- Does the introduction into an adversary system of a software agent with capabilities for both exploitation and destructive action constitute a use of force? It may be relevant to consider as an analogy the insertion into a potential adversary of a human agent skilled both in espionage and in sabotage.
- Does the introduction of a remotely reprogrammable software agent into an adversary system constitute a use of force? A possible analogy in this case may be a preplanted mine that can be detonated by remote control from a long distance away.
- Does a non-destructive probe of an adversary's computer network for intelligence-gathering purposes to support a later cyberattack itself constitute a use of force? (An analogy might be drawn to the act of flying near an adversary's borders without violating its airspace in order to trigger radar coverage and then to gather intelligence on the technical operating characteristics of the adversary's air defense radars. Though such an act might not be regarded as friendly, it almost certainly does not count as a use of force.)

²³ Cited in Roger D. Scott, "Territorially Intrusive Intelligence Collection and International Law," *Air Force Law Review* 46:217-226, 1999; available at [http://permanent.access.gpo.gov/lps28111/Vol.46\(1999\)/scottfx4\[1\].doc](http://permanent.access.gpo.gov/lps28111/Vol.46(1999)/scottfx4[1].doc).

Box 7.7 provides some scenarios in which such questions arise.

7.2.2.2 During Ongoing Hostilities—Applying *Jus in Bello*

If an armed conflict is ongoing, cyberattacks on any military target (e.g., military command and control systems or an adversary’s defense industrial base) would satisfy the condition of military necessity. At the same time, the legality of such use would be subject to the *jus in bello* conditions regarding proportionality, distinction, and so on, just as they would affect decisions involving the use of kinetic weapons in any given instance. Note also that the attack/defense distinction—central to applying *jus ad bellum*—is not relevant in the midst of armed conflict and in the context of *jus in bello*. Some of the issues raised by *jus in bello* for cyberattack are described below.

7.2.2.2.1 Proportionality of Military Action

The proportionality requirement stipulates that military actions be conducted in a way that the military gain likely from an attack outweighs the collateral damage of that attack. For example, the electric power grid is often discussed as a likely target for cyberattack. In a full-scale nationwide mobilization, the electric power grid supports a nation’s war effort, and thus it might appear to constitute valid military targets in a conflict. But for an attack on it to be regarded as proportional, a judgment would have to be made that the harm to the civilian population from disrupting electrical service was not disproportionate to the military advantage that might ensue from attacking the grid.

Several characteristics of cyberattack affect proportionality judgments.

- Predicting and understanding the actual outcome of a cyberattack is very intelligence-intensive—estimates of likely collateral damage and likely intended damage will depend on myriad factors (as discussed in Section 2.3.5). And much of this intelligence will be difficult to obtain, especially on short notice. Thus, the a priori predictions of outcome and actual outcomes will often be highly uncertain. Of course, commanders must proceed even in the face of many uncertainties about the characteristics of the target in both kinetic and cyber targeting, and they are not required to take into account outcomes and effects that are known to be very unlikely. But the open question is how commanders should account for uncertainties in outcome that are significantly greater than those usually associated with kinetic attacks in the sense that there may not be an analytic or experiential basis for estimating uncertainties at all. Under such circumstances, how is the proportionality judgment to be

BOX 7.7 The LOAC Exception for Espionage—Possible Examples

The following examples illustrate possible scenarios that pose questions about whether a given cyber offensive action should be treated as cyberattack or cyberexploitation.

- *A cyber offensive action introduces a two-part software agent into an adversary system.* The software agent is designed with two parts. One part is used for cyberexploitation, monitoring traffic through the system and passing the traffic along to a collection point. A second part is potentially used for cyberattack, awaiting an instruction to “detonate,” at which point it destroys the read-only memory controlling the boot sequence of the machine where it resides. Until the agent detonates, no damage has been caused to the system, and no use of force has occurred. On the other hand, the potential to do damage has been planted, and perhaps the act of planting the agent with a destructive component can be regarded as a threat of force. Under what circumstances, if any, does this offensive action constitute a use of force or the threat of force? The clandestine nature of the agent complicates matters further—an essential dimension of “threat” is that it must be known to the party being threatened, and there is a strong likelihood that the system owner will not know of the agent’s existence. Still, the owner could discover it on its own, and might well feel threatened after that point.

- *A cyber offensive action introduces an upgradeable software agent into an adversary system.* As introduced, the agent is an agent for cyberexploitation, monitoring traffic through the system and passing it along to a collection point. But through a software upgrade transmitted to the agent by clandestine means, the agent can then take destructive action, such as destroying the read-only memory controlling the boot sequence of the machine where it resides. A similar analysis applies in this instance—the agent as introduced does not constitute a use of force, as it has no destructive potential. But it can easily be turned into a destructive agent, and perhaps the act of upgrading the agent with a destructive component can be regarded as a threat of force or an imminent attack. Under what circumstances, if any, does this offensive action constitute a use of force or the threat of force?

- *A probe is launched to map an adversary’s computer network.* As such, this operation is a cyberexploitation—it is gathering intelligence on the network. Such an attack causes no damage to the network but provides the attacker with valuable information that can be used to support a subsequent cyberattack.

made? What is clearly the *wrong* way to account for such uncertainties is to ignore them. Although it is a natural human tendency to ignore factors whose significance is unknown, in practice such behavior amounts to assigning zero weight to them.

- Because the outcome of a cyberattack may depend on very small details known only to the party attacked, such parties may have greater

opportunity to claim collateral damage from a cyberattack when in fact no such damage occurred. And the attacking party might well have a difficult time refuting such claims, even if it were willing to divulge details about the precise nature of the cyberattack in question. So, for example, a cyberattack against an air defense network might lead to claims that the attack also shut down electric power to a hospital. The possibility of false claims exists with kinetic attacks as well, but claims about collateral damage from a cyberattack are likely to be even more difficult to refute.

- The damage assessment of a cyberattack necessarily includes indirect as well as direct effects, just as it does when kinetic weapons are involved. These indirect effects, if they relate to effects on civilians, count in the proportionality judgment. Thus, for example, if a cyberattack to disable a dual-use telephone switching station for several hours is contemplated, the fact that medical patient lives may be lost because the station serves a hospital must be factored into the judgment about whether the attack meets the proportionality requirement if such an outcome can be reasonably foreseen.

- Some cyberattacks are potentially reversible. To the extent that the damage caused is reversible, a lesser amount of collateral damage should also be expected, and thus the calculation of weighing the military utility against collateral damage of a given cyberattack will be tilted more in favor of proceeding rather than refraining from the attack—that is, reversibility will make the action more likely to be proportional, and could result in a cyberattack being preferred to a kinetic attack with all else being equal. (Indeed, even if the military effect is somewhat less, there may still be a LOAC obligation to use the less damaging cyberweapon if the collateral damage would be substantially lower.)

As an example, consider an electric power grid that serves both military and civilian purposes. The grid could be a legitimate military target, even if the civilian use is extensive, as long as the military use is very important to the enemy's war effort. If the grid's control centers are bombed, it may take a very long time to restore service when the war is over, but if they can be shut down by cyberattacks, it may be possible to restore service much more quickly. The military gain is achieved even by a short-term disruption (at least if the cyberattack can be repeated as needed), while in terms of impact on the civilian population there is a big difference between a loss extending for a few weeks or even longer during hostilities and one stretching long into the postconflict reconstruction phase.

7.2.2.2.2 Distinctions Between Military, Civilian, and Dual-Purpose Assets

Under traditional LOAC *jus in bello*, only a nation's military forces are allowed to engage in armed hostilities with another nation. In addition, a nation is entitled to attack combatants but must refrain from attacking non-combatants as long as the latter avoid any participation in the conflict. Cyberattacks raise a number of questions in this context:

- Does compromising the computers of non-combatants violate prohibitions against attacking non-combatants?
- Under what circumstances does a cyberattack on national infrastructure that affects both civilian and military assets constitute a LOAC violation?
- What responsibilities does a nation have to separate civilian and military computer systems and networks?
- Must military computer systems and networks be made identifiable as such to a potential attacker if a nation is to claim immunity for civilian systems and networks?

Box 7.8 provides some possible scenarios in which such questions arise.

7.2.2.2.3 Distinctions Between Military and Civilian Personnel

The LOAC principle of distinction also confers different rights and responsibilities on combatants and non-combatants. Combatants are the only parties who are entitled to use force against the enemy. Combatants must also be trained in the law of war, serve under effective discipline, and be under the command of officers responsible for their conduct.²⁴ Whenever they are engaged in combat operations (and subject to the permissibility of employing a legitimate *ruse de guerre*), they must be identifiable (usually by carrying arms openly and wearing a distinctive uniform) as combatants. Lawful combatants captured by the enemy may not be punished for their combatant acts so long as they complied with the law of war; must be treated in accordance with agreed standards for the treatment of prisoners of war; and must be released promptly at the cessation of hostilities. The enemy is also entitled to target lawful combatants deliberately. Non-combatants have an affirmative duty to

²⁴ Put differently, accountability mechanisms under LOAC are established through the doctrine of superior orders (i.e., someone higher in the chain of command has responsibility for the known or likely actions of someone lower in the chain of command) and the obligation to disobey manifestly illegal orders (someone lower in the chain of command has an obligation to obey lawful orders and a concomitant obligation to refuse to obey orders that are outside the scope of international standards).

BOX 7.8 Ambiguities Raised by Cyberattack Against Dual-Purpose Assets—Possible Examples

The following examples illustrate possible scenarios that raise questions related to attacks on dual-purpose assets.

- *A cyberattack can be routed to its ultimate target through intermediary computers.* If the United States wishes to conduct a cyberattack on Zendia, it may wish to route its attack through the personal computers owned and operated by Zendian citizens. (For example, a botnet used to attack Zendia may well use such computers.) Does the compromise of the Zendian citizen computers constitute an “attack” on Zendian citizens?

One important point is that not all actions that harm the Zendian citizens constitute an attack for LOAC purposes. In the case of a personal computer being compromised for launching a cyberattack against Zendia, the harm to its Zendian owner is minimal, because that computer is likely just as useful to its owner as before. Even if it is not, it is hard to imagine that the owner might die as a result of the compromised computer, or even that the property damage suffered is significant, and, on the assumption that the attack has a proper military objective, any damage to civilian interests would be acceptable as collateral damage.

On the other hand, if the cyberattack was deemed to be a use of force or an armed attack against Zendia, the compromise of Zendian citizen computers to prosecute the attack might be regarded in the same vein—thus making the attacker responsible for attacking civilians. In addition, the Zendian government might well take action against Zendian citizens, with unknown consequences for them (and possibly implicating human rights law, as discussed in Section 7.2.5).

- *A cyberattack can be directed against dual-use assets with both civilian and military uses.* Traditional LOAC allows attacks on dual-use targets if the conditions of military necessity, proportionality, distinction, and discrimination are met. The principle of distinction requires that the attacker distinguish between military and civilian targets and refrain from attacking the latter.

In traditional armed conflict, a combination of visual identification and geography often suffices to identify a valid military target—for example, a tank is easily

refrain from participating in combatant activities, and are legally immune from deliberate targeting;²⁵ non-combatants who engage in combatant activities are subject both to military action and, if captured, to criminal prosecution.

Today, there is a growing dependence of the modern military on

²⁵ Note, however, that the systems used to launch cyberattacks are legitimate military targets, and civilians who qualify for the narrow category of “civilians accompanying the armed forces” (presumably those who operate and maintain those systems)—even if they do not actually press the button that launches a cyberattack—are both eligible for prisoner-of-war status and also legitimate military targets for the enemy.

recognizable as a military vehicle and, if it is located behind enemy lines, can reasonably be presumed to be an enemy vehicle. But a computer is not so easily recognized, as both its functionality and geographic location are often not easily available to a would-be attacker.

For example, the commingling of civilian and military communications channels on media such as the Internet or the public switched telephone network might provide an adversary with a plausible military rationale for attacking facilities associated with these media. Moreover, given that civilian and military computer systems can be difficult to distinguish, a question arises as to whether a nation that does not provide machine-readable indications of a computer's status (military or civilian) would have the right to challenge the legality of a cyberattack that damaged or destroyed a civilian computer.

- *A large-scale cyberattack can be directed against (elements of) the critical infrastructure of a nation.* As noted earlier, restraints on the use of biological and chemical weapons exist in part because they are inherently non-discriminating weapons. Although there is no specific ban on the use of non-discriminating weapons per se, the proportionality requirement means that the military value of a given attack must be weighed against collateral damage. LOAC requires military forces to refrain from using a non-discriminating weapon when a more discriminating weapon would be equally effective, and also to refrain from attacking a military target when the only available means to do so is likely to cause disproportionate civilian damage.

In a cyberattack context, this prohibition appears likely to apply to attacks that cannot be limited to specific (military) targets. Thus, a computer network attack based on the Morris worm, for example, might be prohibited, because its effects were wholly indiscriminate and no effort was made to discriminate between appropriate and inappropriate targets. The open question is whether the harm caused to civilians rises to a level that qualifies as disproportionate. Mere inconveniences would not, but death on a large scale would. In between are cases such as the inability to conduct financial transactions electronically, periodic interruptions in electrical power, major disruptions in travel and transportation schedules, and outages in communications capability.

civilians and civilian-provided services and expertise that blurs traditional distinctions between military and civilian activity and personnel. As a legal matter, civilians formally attached to the armed forces (e.g., as contractors) are entitled to some of the privileges of combatants (such as prisoner-of-war status if captured). Civilians engaged in self-help activities (which might resemble combatant activities) are subject to the regular criminal laws.

In light of the often-specialized expertise needed to launch computer network attacks (expertise that may be provided by civilians), an important question is thus raised about what it means to "launch" an attack or

BOX 7.9 Drawing the LOAC Line for Civilian Immunity—Possible Examples

In a war involving the United States, civilians working in a U.S. munitions plant are likely not to enjoy LOAC protection from attack, as they are making a direct contribution to the U.S. war effort. In a cyber context, one can imagine several gradations of civilian involvement in launching a cyberattack, and where the line of LOAC protection should be drawn is an open question. That is, in which of the following scenarios is the civilian entitled to LOAC protection?

- A civilian posts a vulnerability notice for the open source Linux operating system that a U.S. cyberattack exploits.
- A civilian contractor for the DOD identifies the presence of this vulnerability on a Zendian system.
- A civilian exploits the vulnerability by introducing a hostile agent into the Zendian system that does not damage it but that can be directed to cause damage at a subsequent time.
- A civilian dictates to a military officer the precise set of commands needed to activate the hostile agent.

to “use force against an enemy.”²⁶ Box 7.9 describes a possible continuum of civilian involvement in cyberattack.

In addition, the instruments of cyberattack—cyberweapons—are easily available to private groups and individuals as well as governments, thus raising the possibility that private groups and individuals could join a conflict nominally prosecuted between nation-states. This point is further discussed in Section 7.2.3.3 below.

7.2.2.2.4 Neutrality in a Cyberattack

A cyberattack that is conducted at a distance, especially one conducted over the Internet, is likely to involve message traffic that physically transits a number of different nations. A cyberattack on Zendia initiated by the U.S. government may first be transmitted to Ruritania and then to Armpitia and finally to Zendia. Moreover, it is entirely possible, likely even, that neither Ruritania nor Armpitia would be aware of the fact that they were carrying attack traffic at all.

²⁶Such a question applies in many other contexts, such as civilians flying missile-armed drones remotely, designing nuclear weapons, or working in an ammunition or uniform-making factory.

- Given that a computer of military significance can be located anywhere in the world, under what circumstances—if any—is it entitled to protection under LOAC provisions for neutrality?
 - What, if any, are the obligations of neutral nations to prevent cyberattacks from emanating from their territory? (“Emanating from X” means that X is an intermediate node in the attack pathway.)
 - What, if any, are the obligations of belligerents to avoid routing cyberattacks through the computers of neutral nations?

Box 7.10 provides some scenarios in which such questions arise.

A paper by George Walker addresses some of the issues that arise in scenarios similar to those described in Box 7.10.²⁷ Walker notes that legal guidance regarding information warfare (cyberattack) and neutrality will have to be found by analogy to existing international law, since existing law on neutrality does not address issues related to cyberwarfare. He argues that some LOAC principles, such as those related to telegraphy, will apply to Internet messages and more conventional communications, and further that there are many principles—primarily in the law of naval warfare but also some from the law of land and air warfare—that may be cited by analogy in cyberwarfare involving neutrals. His reasoning is based on the premise that aerial warfare and especially naval warfare are conducted in “fluid” mediums, much like the Internet’s electronic pathways that are, like the high seas, no nation’s property. He also points to a relatively well-developed set of rules or general principles in the law of the sea, the law of naval warfare, and the law of air warfare, from which useful analogies for information warfare may be drawn. As an example of a useful analogy from the law of naval warfare, Walker suggests that a nation aggrieved by cyberattacks should have the right to take such actions as are necessary in the territory of a neutral that is unable (or perhaps unwilling) to counter enemy cyberattacks making unlawful use of that territory.

A contrary conclusion might be drawn by an analogy to telephone and telegraph communications as they were handled in the Hague Convention of 1907. Section 5, Article 8 of that convention stipulates that a neutral nation need not “forbid or restrict the use on behalf of the belligerents of telegraph or telephone cables or of wireless telegraphy apparatus belonging to it or to companies or private individuals.”²⁸ If there is no obligation of the neutral to stop the transit, there is no right of the belligerents to act against the transit. If the analogy between telegraph/telephone communications

²⁷ George K. Walker, “Information Warfare and Neutrality,” *Vanderbilt Journal of Transnational Law* 33(5):1079-1200, November 2000.

²⁸ See http://avalon.law.yale.edu/20th_century/hague05.asp#art5.

BOX 7.10 Cyberattack and Neutrality—Possible Examples

The following examples illustrate possible scenarios that raise questions related to neutrality in a cyberattack context.

- *A cyberattack is launched against Zendia that requires transit through Ruritania, a declared neutral nation.* In this instance, the open question is whether the transiting of a cyberattack is more like an overflight by military airplanes (in which case the intermediate nation has an obligation to stop such overflights or allow the other belligerent to do so) or more like the use of telephone and telegraph cables that are provided impartially to both sides (in which case the 1907 Hague Convention explicitly states that the intermediate nation is *not* obligated to prevent such use).¹

- *During conflict between the United States and Zendia, a U.S. cyberattack is launched on a computer controlling production in a Zendian defense plant. However, the computer itself is located in Ruritania, a declared neutral nation that provides computerized production control services to any nation willing to pay for them.* A question arises because the effects of attacking a given computer may not be felt at all in the immediate geographic vicinity of the computer, thus raising the question of which geographic location is relevant to the determination of legitimacy for attack. That is, is the computer operating in Ruritania a valid military target?

- *A cyberattack is launched against the United States by an unknown party that depends on the use of compromised computers belonging to citizens and companies of Ruritania, a declared neutral nation.* Under the doctrine of “self-defense in neutral territory,” Ruritania must take action that eliminates the threat (in this case, the cyberattack) emanating from its territory,² allow or assist the United States to do so itself, or possibly face the consequences of a response from the United States. Complications arise regarding the sequencing of a self-defense response, because in the time that it takes to make a determination that Ruritania is unwilling or unable to stop the cyberattack, the damage to the United States may have been done, or the opportunity for an effective self-defense response lost.

¹ Note an interesting side effect of a policy decision to avoid routing through neutral nations. If U.S. policy required avoidance of routing through neutrals, and if a target nation knew that policy, then said target could effectively shield itself from U.S. cyber operations by peering only with neutrals.

² At the very least, such action would require the government of the putatively neutral nation to have the legal standing to stop such behavior and to demonstrate some plausible degree of cooperation in doing so.

and packet-switched Internet communications is valid (in both cases, the country transited has no real way of knowing the ultimate destination of transiting messages, and selective interference with the communications of belligerents is not practical), an analyst might conclude that belligerents do not have the right to interfere with nodes located in the neutral nation.

7.2.2.2.5 *Covert Action*

Covert action is statutorily defined in the United States as “an activity or activities of the United States Government to influence political, economic, or military conditions abroad, where it is intended that the role of the United States Government will not be apparent or acknowledged publicly.”²⁹ As discussed in Section 7.3.1, U.S. domestic law addresses agency responsibilities within the U.S. government—covert action is the responsibility of the intelligence agencies, whereas military activities are the responsibility of the Department of Defense.

At the same time, international law is not sensitive to which agencies of a given government take action. This fact has at least two implications. First, *jus ad bellum* and the UN Charter apply to covert action—and an action with a scale of effect that would constitute a use of force or an armed attack if performed by U.S. military forces would be regarded in the same way even if it were designated as covert action by the President of the United States. Second, *jus in bello* would apply to any U.S. covert action involving the use of cyberattack during armed conflict.

7.2.2.2.6 *An Operational Note—Jus in Bello in Practice*

U.S. military commanders undergo formal training in the laws of armed conflict so that they can appropriately direct their forces during combat. In most cases, senior commanders have the assistance of lawyers who can and do review a proposed course of action (such as an air tasking order) for LOAC compliance.

Operating under combat conditions, commanders with significant experience in a particular kind of situation and with particular weapons have a good intuition for the outcome of a legal review of a proposed course of action. A LOAC review may result in adjusting the parameters of an attack at the margins, but the outcome of the review is largely a given (that is, the proposed action will be allowed) because the commander with a certain amount of accumulated experience is unlikely to propose a course of action that is far outside the boundaries of what a legal review would allow. Under such circumstances (that is, in a kinetic war), the LOAC review process can be expedited if and when the commander and the lawyers have both internalized the same general outline of what is and is not allowable under their shared legal paradigm.

But when there is little or no experience on which to draw, the congruence between the course of action proposed by commanders and what

²⁹ 50 USC 413b(e). See also Joint Explanatory Statement of the Committee of Conference, H.R. 1455, July 25, 1991, Intelligence Authorization Act of FY 1991.

the lawyers would say is more likely to break down. This is particularly relevant if cyberweapons—which have been used much less often in combat compared to kinetic weapons—are to be used. Today, relatively few commanders have substantial experience with cyberattack, and relatively few military lawyers have experience in rendering LOAC judgments about cyberattack (and lawyers are often reluctant to set new precedents in practice). Thus, where cyberattack is concerned, it is less likely that commanders and lawyers will have internalized similar boundaries of what is and is not acceptable.

One important consequence of this state of affairs is that one might expect LOAC review of cyberattack plans to be more challenging than review of kinetic attack plans. Consistent with this point, James Miller, former deputy assistant secretary of defense for requirements, plans and counterproliferation reported to the committee that because of the potential for unintended effects in cyber networks and sensitivity to the vulnerability of U.S. networks as well as the precedent-setting nature of decisions, LOAC review for cyber operations in Kosovo was indeed very challenging.

7.2.2.3 A Summary of Applying LOAC to Cyberattack

During acknowledged armed conflict (notably when kinetic and other means are also being used against the same target nation), cyberattack is governed by all the standard LOAC criteria of *jus in bello*—military necessity, proportionality, distinction, and so on, although the legal analysis in any given situation involving cyberattack may be more uncertain because of its novelty relative to kinetic weapons.

In other cases (that is, in less than acknowledged armed conflict), the legal status of a cyberattack is judged primarily by its effects, regardless of the means, according to the criteria of *jus ad bellum* and of the Charter of the United Nations. Therefore, if the effects (including both direct and indirect effects) to be produced by a cyberattack would, if produced by other means, constitute an armed attack in the sense of Article 51 of the UN Charter, it is likely that such a cyberattack would be treated as an armed attack. Similarly, if a cyberattack had the same effects and was otherwise similar to governmentally initiated coercive/harmful actions that are traditionally and generally not treated as the “use of force” (e.g., economic sanctions, espionage, or covert actions such as planting information or influencing elections), such a cyberattack would likely not be regarded as an action justifying a use of force in response.

7.2.3 International Law and Non-state Actors

International law binds nations, and only in exceptional cases binds non-state actors such as corporations, individuals, or terrorist groups. However, there are both domestic and international legal doctrines that restrict, and in most cases prohibit, non-state actors from actions that would be international use of force if undertaken by nation-states, and nations do have obligations in some circumstances to prevent these actors from acting in such ways that violate international law.³⁰

7.2.3.1 International Law and Non-state Actors—Terrorists

Traditional LOAC emerged from the need to regulate nation-to-nation conflict between national military forces. But other forms of conflict in the 1990s and 2000s (such as terrorism) have blurred many of the distinctions between the LOAC and domestic law enforcement.

In such instances, both military and civilian dimensions are relevant and raise questions about the applicability of LOAC and law enforcement approaches.³¹ The difficulties arising are hard enough to resolve when the aggressive act is a tangible action—that is, the use of deadly force to harm persons or destroy property. But they are compounded when the aggressive act is in cyberspace and its harm can only be assessed by consequences that are not fully knowable except with the passage of time.

These issues come to the fore in an international security environment involving subnational groups and non-state actors. In this new

³⁰ Even prior to the September 11, 2001, attacks on the United States, a nation-state was responsible for the acts of private groups inside its territory over which it exercised “effective control.” (See, for example, Article 8 of the ILC (International Law Commission) State Responsibility Articles, available at http://untreaty.un.org/ilc/texts/instruments/english/commentaries/9_6_2001.pdf, pp. 47 ff; and the ICJ (International Court of Justice) Nicaragua decision (arguing for “effective control”) and the ICTY (International Criminal Tribunal for Yugoslavia) Tadic decision (arguing for “overall control”).) In the aftermath of those attacks, the United States took the position that the mere harboring of these actors, even in the absence of control over them, suffices to make the state where the terrorists are located responsible for their actions (UN Security Council, “Letter Dated 7 October 2001 From the Permanent Representative of the United States of America to the United Nations Addressed to the President of the Security Council,” UN Doc. No. S/2001/946 (2001)), and many parts of the international community, including the UN Security Council, concurred with this position (see Derek Jinks, “State Responsibility for the Acts of Private Armed Groups,” *Chicago Journal of International Law* 4(1):83-96, Spring 2003).

³¹ Human rights advocates sometimes assert that human rights law also applies even when LOAC applies. Although this assertion is categorically rejected by the U.S. government, the political reality is that this argument is likely to resonate with some outside observers and thereby raise the level of world scrutiny for all U.S. uses of military force, thus adding to the political pressures on the United States in a crisis.

environment, questions have arisen as to whether terrorists are subject to LOAC, to criminal law, or to some other body of law that has yet to be established. Although the Supreme Court held that common Article 3 of the 1949 Geneva Conventions applied to the conflict against Al Qaeda and the Taliban authorized by Congress on September 14, 2001 (*Hamdan v. Rumsfeld*), it is generally fair to say that the details of how LOAC does and does not apply in a conflict with terrorists are far less developed and clear than in a conflict between nation-states.

Nevertheless, a number of practical considerations arise in dealing with non-state actors given that these actors (call them terrorists for now) will almost surely be operating from the territory of some nation-state. Therefore, any action taken against them may raise issues about violating the sovereignty of that nation and its rights and obligations with respect to terrorist operations from or through its territory.

All of the above issues apply in contemplating cyberattacks as they might be conducted by terrorists. Cyberattack weapons are inexpensive and easily available but may have the potential to cause widespread damage and destruction, characteristics that may make such weapons attractive to terrorists.

The important question is whether, when, and why a cyberattack by a non-state actor should be treated primarily as a law enforcement matter, a national security matter, or a mix of the two. (The first manifestations of a cyberattack are likely to require investigation to determine its source. But once such a determination has been made, this threshold question will inevitably arise.)

One relevant question in making such a determination is whether the attack has serious enough consequences (death or destruction) that it would qualify as a use of force or an armed attack on the United States had it been carried out with kinetic means. A second question concerns the geographic origin of the attack. A third question may be the nature of the party responsible for the attack (e.g., national government, terrorist group). As a factual matter, none of these pieces of information may be known at the time the attack becomes known (as discussed in Section 2.4.1 on tactical warning and attack assessment); nevertheless, these questions will be prominent in the minds of senior decision makers because the answers may have profound implications for the legitimacy of a response.

If and when the geographic origin of the attack becomes known (call it *Zendia*), *Zendia* may have one of several stances toward cooperation with the United States. At one extreme, *Zendia* may cooperate fully with the United States in stopping the attack emanating from its soil, where full cooperation can mean anything from placing *Zendian* law enforcement

and security services at U.S. disposal to giving permission for the United States to act as it sees fit in its response to the attack. At the other extreme, Zendia may simply refuse outright any and all U.S. requests for assistance and cooperation. And Zendia's cooperation may fall onto *any* point along this spectrum, raising a variety of legal and policy issues. For example:

- Even if Zendia wishes to cooperate fully, it may not have the legal authority to address the hostile activity in question. That is, the activity may not violate any Zendian law. If Zendia is a signatory to the Cyber-crime Convention, it is obligated to extend such cooperation if the cyber activity emanating from Zendian soil is considered a criminal matter under Zendian law. Nevertheless, not all nations are signatories to the convention, and the convention itself is oriented toward a law enforcement approach (that is, investigation, arrest, prosecution, and legal due process) that is often too slow given how rapidly a cyberattack can unfold. Finally, "permission" can be ambiguous, as in those instances when there is some doubt or question about who speaks for the "legitimate" government of Zendia.

- If Zendia explicitly refuses to cooperate, the United States could assert the right of self-defense in neutral territory discussed in Section 7.2.1.2. To be sure, such a decision would be a policy decision and would depend on a host of factors such the scope and nature of the proposed U.S. action, whether Zendia is capable of resisting unilateral U.S. actions taken in response, and other areas of U.S.-Zendian cooperation or contention. (For example, if Zendia has nuclear weapons capable of reaching U.S. targets, the decision-making calculus for policy may change considerably though the legal issues do not.)

- Perhaps the most problematic response is a posture of limited, grudging, or excessively slow Zendian cooperation, or words that indicate cooperation but are unaccompanied by matching actions. For example, permission for the United States to undertake various actions might be slow in being granted, or unduly circumscribed in a way that impeded further investigation or action; information provided to the United States might be incomplete. Under these circumstances, the Zendian response could conceivably take a very long time and would be unlikely to be fully satisfactory to the United States. Yet even if the response is inadequate for U.S. purposes, it might still be enough to sway the court of world opinion against an aggressive U.S. response and perhaps even to forestall it. Even though a deliberate stalling is probably equivalent to an outright refusal to cooperate, making the determination that Zendia is being deliberately uncooperative may be problematic in the absence of an explicit statement.

7.2.3.2 International Law and Non-state Actors—Multinational Corporations

Barkham³² notes that many multinational corporations exercise power and influence that at times rivals those of small nation-states. International law (including LOAC) also does not directly constrain the actions of such corporations to any significant extent. On the other hand, they are subject to the laws of those nations in which they have a presence, and sometimes those laws result from government-to-government agreements (of which the Convention on Cybercrime (discussed below in Section 7.2.4) is an example).

Of significance to this report is the fact that certain multinational corporations will have both expertise and resources to launch cyberattacks of a significant scale should they choose to do so. If they did, such multinational corporations might threaten cyberattacks against weak nation-states to gain concessions or launch cyberattacks against economic competitors to place them at a competitive disadvantage (e.g., by disrupting production).

7.2.3.3 International Law and Non-state Actors—Patriotic Hackers

LOAC presumes that armed conflict is initiated only at the direction of government, only by its authorized military agents, and specifically not by private groups or individuals. Thus, governments maintain armed forces to participate in armed conflict, under the government's direction.

But in the Internet era, another type of non-state actor that complicates the legal landscape for cyberattack is the "hactivist" or patriotic hacker. During times of conflict (or even tension) with another nation, some members of a nation's citizenry may be motivated to support their country's war effort or political stance by taking direct action (Box 7.11). Hactivists or patriotic hackers are private citizens with some skills in the use of cyberattack weapons, and they may well launch cyberattacks on the adversary nation on their own initiative, that is, without the blessing and not under the direction or control of the government of their nation.

Apart from their possible operational interference with other, government-authorized actions, the actions of these patriotic hackers may greatly complicate the conduct of diplomatic action. For example, if Zendian patriotic hackers launch cyberattacks against the United States, the United States is entitled to respond as though the Zendian government were

³² Jason Barkham, "Information Warfare and International Law on the Use of Force," *New York University International Law and Politics* 34:57-113, 2001.

responsible. Whether it should do so is a policy question.³³ What if the patriotic hackers are part of the Zendian diaspora and are located in territories other than Zendia? What actions should the United States take to respond to Zendian patriotic hackers if the Zendian government says in response to a U.S. inquiry, “We do not endorse or encourage these attacks by our citizens, but at the same time, they are not doing anything that we have the ability (or perhaps the legal authority) to stop, so the best thing for you to do is to cease your aggressive actions against Zendia.”? Note the similarity of this situation involving Zendian patriotic hackers to the situation discussed in Section 7.2.3.1 involving cyberterrorists operating from Zendian soil.

As noted in Section 7.2.1.2, states likely have affirmative obligations to refrain from harboring perpetrators of terrorist attacks. To the extent that Zendia supports patriotic hackers in their activities, other nations targeted by these parties may have a legitimate complaint to bring forward in an appropriate international tribunal by asserting that Zendia is indeed harboring perpetrators of terrorist activity—indeed, these other nations may well be entitled to invoke inherent rights of self-defense consistent with Article 51. One significant question in this regard is whether a failure to suppress the activities of patriotic hackers should count as support for them.

7.2.4 The Convention on Cybercrime

The Convention on Cybercrime commits signatories to the adoption of “a common criminal policy aimed at the protection of society against cybercrime . . . by adopting appropriate legislation and fostering international co-operation.”³⁴ The convention establishes a common minimum standard of relevant offenses to be applied at the national level in several areas. Five criminal offenses are specifically defined to protect the confidentiality, integrity, and availability of computer data and systems, namely:

- *Illegal access*—intentional access to the whole or any part of a computer system without right, where the offence may be considered to have

³³ As a precedent, the International Court of Justice held in the 1980 *U.S. v. Iran* case that the actions of a state’s citizens can be attributed to the government if the citizens “acted on behalf on [sic] the State, having been charged by some competent organ of the Iranian State to carry out a specific operation.” Further, the court found that the Iranian government was responsible because it was aware of its obligations under [international law] to protect the U.S. embassy and its staff, was aware of the embassy’s need for help, had the means to assist the embassy, and failed to comply with its obligations. See United States Diplomatic and Consular Staff in Tehran (*U.S. v. Iran*), 1980 I.C.J. 3, 29 (May 24). Cited in Barkham, 2001.

³⁴ Council of Europe, Convention on Cybercrime, November 23, 2001.

BOX 7.11 Hacktivism During International Conflict and Tension

A number of incidents of privately undertaken cyberattacks have been publicized:

- Immediately after the start of the second intifada in Israel in late September 2000, Palestinian and Israeli hackers conducted a variety of cyberattacks on each other's national web presences on the Internet.¹
- In the aftermath of the early 2001 incident between the United States and China in which a U.S. EP-3 reconnaissance aircraft collided with a Chinese F-8 interceptor, both Chinese and U.S. hackers attacked the web presence of the other nation. In both cases, attacks were aimed mostly at website defacement and denial of service.²
- In the wake of the May 1999 bombing by the United States of the Chinese embassy in Belgrade, the U.S. National Infrastructure Protection Center issued an advisory (NIPC Advisory 99-007) noting "multiple reports of recent hacking and cyber activity directed at U.S. government computer networks, in response to the accidental bombing of the Chinese embassy in Belgrade. . . . Reported activity include[d] replacing official web pages with protest material and offensive language, posting similar language in chat rooms and news groups, and denial of service e-mail attacks."³
- American hackers have been known to attack jihadist websites. For example, an American was reported by *Wired* to have hijacked www.alneda.com, a widely used website for jihadist recruitment.⁴ His motive for doing so was said to be a decision made after the September 11 attacks: "I was going to use every skill I had to screw up the terrorists' communication in any way I could."

occurred if security measures are infringed with the intent of obtaining computer data or other dishonest intent or where computer systems are networked.

- *Illegal interception*—intentional interception without right, made by technical means, of non-public transmissions of computer data to, from, or within a computer system, including electromagnetic emissions from a computer system carrying such computer data.

- *Data interference*—intentional damage, deletion, deterioration, alteration, or suppression of computer data without right.

- *System interference*—intentional serious hindering without right of the functioning of a computer system by inputting, transmitting, damaging, deleting, deteriorating, altering, or suppressing computer data.

- *Misuse of devices*—intentional production, sale, procurement for use, import, distribution, or possession of a computer password, access code, or device, including a computer program, designed or adapted primarily for the purpose of committing any of the other four offenses.

- Russian hackers are widely reported to have been responsible for the cyberattacks on Estonia in 2007 (see Box 3.4 in Chapter 3) and Georgia in 2008.⁵

Allen and Demchak generalize from experiences such as these to predict that future conflicts between nations may involve:

- Spontaneous attack action in cyberspace by “patriots” on each side.
- Rapid escalation of their actions to a broad range of targets on the other side. Allen and Demchak posit that because “hacktivists” are interested in making a statement, they will simply attack sites until they find vulnerable ones.
- Involvement of sympathetic individuals from other nations supporting the primary antagonists.

¹ Associated Press, “Cyberwar Also Rages in Mideast,” October 26, 2000, available at <http://www.wired.com/politics/law/news/2000/10/39766>.

² Michelle Delio, “A Chinese Call to Hack U.S.,” *Wired*, April 11, 2001, available at <http://www.wired.com/news/politics/0,1283,42982,00.html>.

³ See NIPC Advisory 99-007, available at <http://www.merit.edu/mail.archives/netsec/1999-05/msg00013.html>.

⁴ Patrick Di Justo, “How Al-Qaida Site Was Hijacked,” *Wired*, August 10, 2002, available at <http://www.wired.com/culture/lifestyle/news/2002/08/54455>.

⁵ “Expert: Cyber-attacks on Georgia Websites Tied to Mob, Russian Government,” *Los Angeles Times*, August 13, 2008, available at <http://latimesblogs.latimes.com/technology/2008/08/experts-debate.html>.

SOURCE: Adapted largely from Patrick D. Allen and Chris C. Demchak, “The Palestinian-Israeli: Cyberwar,” *Military Review* 83(2), March-April 2003.

Criminal possession may be defined as the possession of a number of such devices. No criminal liability is imposed where the intent is for reasons other than to commit any of the other four offenses.

The Convention on Cybercrime also identifies a number of ordinary crimes that are often committed through the use of computer systems, including forgery and fraud.

The convention defines a computer system to be “any device or a group of interconnected or related devices, one or more of which, pursuant to a program, performs automatic processing of data.” Computer data is defined to be “any representation of facts, information or concepts in a form suitable for processing in a computer system, including a program suitable to cause a computer system to perform a function.”

The convention calls for signatories to adopt domestic laws that criminalize the above offenses, to provide domestic law enforcement agencies with the authorities and powers necessary for the investigation and

prosecution of such offenses (as well as other offenses committed using a computer system), and to establish an effective regime of international cooperation, including provisions for extradition and mutual law enforcement assistance. Notably, the convention does not establish espionage as an act that violates international law.

As of December 21, 2007, 43 nations had signed the Convention on Cybercrime, of which 21 have ratified it.³⁵ The U.S. Senate ratified the convention in August 2006 and took the view that prior U.S. legislation provided for all that the convention required of the United States. Many but not all European nations have also ratified the treaty.

The convention is significant to the extent that it commits the parties to regard the commission of the various listed offenses as matters that are actionable for the law enforcement authorities of the nation in whose jurisdiction the offenses were committed. (The convention is silent on what actions may be taken by the nation of the victim of such offenses.) That is, if these offenses are committed within the jurisdiction of a signatory nation, that nation is obligated to respond to them as criminal acts—and in particular is required to establish mechanisms for law enforcement cooperation to investigate and prosecute these acts should they occur. Thus, if a cyberattack is launched on the United States in a way that involves another signatory to the Convention on Cybercrime, that nation is obligated to cooperate with the United States in trying to identify the perpetrator.

Of course, not all of the nations of the world have signed on to the Convention on Cybercrime, and a nation's prosecution of cybercriminals and/or its cooperation with an attacked state may be less than zealous. Indeed, the convention also allows a signatory to refuse to provide assistance if the request for assistance concerns an offense that the signatory considers a political offense or if carrying out the request is likely to prejudice the signatory's sovereignty, security, public order, or other essential interests.³⁶ The signatory may also postpone action on a request if such action would prejudice criminal investigations or proceedings conducted by its authorities.³⁷

In short, a signatory nation may decline to cooperate with its obligations under the convention on fairly broad grounds, and the convention lacks an enforcement mechanism to assure that signatories will indeed cooperate in accordance with their obligations. Even in the case of a fully cooperative nation, it may still take a long time to identify a perpetrator

³⁵ See <http://conventions.coe.int/Treaty/Commun/ChercheSig.asp?NT=185&CM=8&DF=&CL=ENG>.

³⁶ Council of Europe, Convention on Cybercrime, Article 27(4).

³⁷ Council of Europe, Convention on Cybercrime, Article 27(5).

and to use legal means to shut down his/her/its criminal cyber activity. Thus, the Convention on Cybercrime would appear to have limited utility in addressing hostile cyberattacks on a prompt time scale, and none at all if a nation refuses to cooperate on any of the broad grounds described above.

7.2.5 Human Rights Law

Human rights are restraints on the actions of governments with respect to the people under their jurisdiction. They can be national in origin (i.e., the civil and political rights under the U.S. Constitution), they may be contained in an international human rights treaty (i.e., the Convention on the Elimination of Discrimination Against Women), or they may be inherent in customary international law.

A central point of contention in human rights law today is the extent of its applicability in situations in which the law of armed conflict is operative, that is, in acknowledged armed conflict or hostilities. The position of the U.S. government is that the moral and ethical imperatives of minimizing unnecessary human suffering are met by the requirements of LOAC (*jus in bello*), and thus that human rights law should not place additional constraints on the actions of its armed forces. By contrast, many human rights observers and non-government organizations would argue that human rights law can and should apply as well as LOAC (*jus in bello*) in acknowledged armed conflict.

As for the governing regime prior to armed conflict, the relevant question is the extent to which human rights law applies before the considerations of *jus ad bellum* are addressed, that is, before combat.

The major treaty relevant to human rights law is the International Covenant on Civil and Political Rights (ICCPR), ratified by the United States in September 1992. Although a variety of human rights organizations strongly disagree, the United States has argued that the ICCPR does not apply extraterritorially, and so it would not regulate U.S. behavior in other countries. This position is based on the text of Article 2 of the ICCPR (the Covenant applies to “. . . all individuals within its territory and subject to its jurisdiction . . .”) and supported by the negotiating history.

If the U.S. position is accepted, cyberattacks that do not rise to the level of armed conflict have no implications from an ICCPR/human rights perspective. If the contrary position is accepted, then two of the rights enumerated in the ICCPR may be relevant to the cyber domain in particular. Article 17 (protecting privacy and reputation) might speak to cyberattacks intended to harm the reputation of an individual, e.g., by falsifying computer-based records about transactions in which he or she had engaged, or to uncover private information about an individual.

Article 19 (protecting rights to seek information) might speak to cyberattacks intended to prevent citizens from obtaining access to the Internet or other telecommunications media.

A variety of other rights, such as the right to life, are potentially relevant as well, although they do not seem as closely tied to the cyber domain. Respecting these other rights would suggest, for example, that a cyberattack intended to enforce economic sanctions would still have to allow transactions related to the acquisition of food and medicine.

7.2.6 Reciprocity

Although U.S. policy will be based on an analysis of what future legal regime would best serve the interests of the United States (including whatever political value can be found in asserting the stance), that analysis must take into account the extent and nature of the effects of such regimes on other parties, both other nation-states and subnational entities, and the likelihood that these other parties might feel obligated to comply with such a regime.

For example, the United States may decide that an expansive definition of “use of force” prohibiting most uses of cyberattack would help to protect the viability of the U.S. information technology infrastructure in the face of international threats. But such a definition would also prohibit most prekinetic conflict uses of cyberattack by the United States as well. Alternatively, it may decide that other key nations would not comply with an expansive definition,³⁸ and thus that a restrictive definition might better serve U.S. interests by allowing most uses of cyberattack.

7.3 DOMESTIC LAW

As noted in Section 7.1, domestic law (which includes the Constitution of the United States, federal statutes, and self-executing treaties) constrains both government institutions and private individuals. For example, U.S. domestic law regulates the division of labor regarding operational activities between the DOD and the intelligence agencies for reasons of government accountability and oversight. Generally, activities of the Department of Defense (DOD) are governed by Title 10 of the U.S. Code, and activities of the intelligence community (IC) by several sections of Title 50. U.S. domestic law also provides substantive law governing

³⁸ Many analysts believe that China is an example of a nation that might well be unwilling to give up a cyberattack-based avenue of asymmetric confrontation against the United States. See for example, Timothy Thomas, *Decoding the Virtual Dragon*, Foreign Military Studies Office, Fort Leavenworth, Kans., 2007.

what private parties can and cannot do, both through highly cyber-specific statutes and more general laws on property, self-defense, and so on.

In general, a state is entitled to use any method for law enforcement within its territory or with respect to its citizens that is consistent with its domestic law. Within the United States, domestic law regulates police conduct and electronic surveillance, and imposes limits on searches or arrests without probable cause and on the unreasonable use of force in making lawful arrests or during other enforcement activities. Under international law, a state must avoid conduct that amounts to torture, genocide, or other blatant and generalized violations of human rights described in the ICCPR.

7.3.1 Covert Action and Military Activity

Chapter 4 addresses some of the operational and policy considerations underlying covert action. But the legal framework governing covert action is also important.

As noted in Chapter 4, covert action has a statutory definition. However, the 1991 Intelligence Authorization Act also included a provision, now codified at 50 USC 413b, that distinguished between covert actions and “traditional military activities,” “traditional counterintelligence activities,” “traditional diplomatic activities,” and “traditional law enforcement activities.” The legislation does not define any of the traditional activities, but the conference report stated the intent of the conferees that:³⁹

“traditional military activities” include activities by military personnel under the direction and control of a United States military commander (whether or not the U.S. sponsorship of such activities is apparent or later to be acknowledged) preceding and related to hostilities which are either anticipated (meaning approval has been given by the National Command Authorities for the activities and for operational planning for hostilities) to involve U.S. military forces, or where such hostilities involving United States military forces are ongoing, and, where the fact of the U.S. role in the overall operation is apparent or to be acknowledged publicly. In this regard, the conferees intend to draw a line between activities that are and are not under the direction and control of the military commander. Activities that are not under the direction and control of a military commander should not be considered as “traditional military activities.”

Covert action requires a written presidential finding in advance of the action that the action is necessary to support identifiable foreign policy

³⁹ Conference Report on H.R. 1455 (House of Representatives), July 25, 1991, available at http://www.fas.org/irp/congress/1991_cr/h910725-ia.htm.

objectives of the United States, submission of the finding to the chairmen of the congressional intelligence oversight committees, and notification of congressional leaders of the action. By contrast, no findings, special approval, or notification are needed for conducting any of the traditional military activities, although activities conducted by the uniformed military are subject to the guidance of and restrictions imposed by the law of armed conflict, and, in practice, many highly sensitive military operations—if conducted outside the framework of a general armed conflict—have been brought to the attention of congressional leadership.

Finally, 50 USC 413b(f) states that “no covert action may be conducted which is intended to influence United States political processes, public opinion, policies, or media.” In practice, U.S. decision makers have sometimes interpreted this provision to mean that no covert action may be conducted that is likely to have such an effect in the United States. Under this interpretation, the use of cyberattack to disseminate false information as part of a covert action might be illegal if such information made it back to the U.S. news media.

The matter is complicated by the fact that for certain kinds of covert action, DOD assets will be needed to execute the applicable plans. Under such circumstances, it is less clear whether the planned action is or is not subject to notification as covert action. In addition, because the mechanism for covert action authorization calls generally for the notification of the appropriate congressional leaders, delay in execution may be possible and negotiation about its terms may be necessary if these leaders object to the action.

The domestic legal requirements for undertaking a covert action require only that the President personally find that the action supports identifiable foreign policy objectives of the United States and that the action is important to the national security of the United States. Thus, as a legal matter, the requirements for a finding regarding an action employing lethal force are the same as for a finding not employing lethal force, and a covert action may use enough lethal force (or destructive force) that it would clearly be a “use of force,” where “use of force” is used in the sense of the UN Charter. Nevertheless, as a practical matter, congressional overseers and executive branch managers of covert actions are more likely to pay more attention to actions that result (or could result) in death and destruction than those that do not. The same is true for covert actions that are likely to be disclosed, or likely to result in failure, or in friendly personnel being captured.

Given this legal environment, it is not surprising that executive branch decision makers have adopted an expansive view of actions that might be considered traditional military activities, and that includes actions that have a very direct military effect on potential military adversaries—even

if such actions would constitute covert action if undertaken by the intelligence community. Indeed, in recent years (that is, since the terrorist attacks of September 11, 2001), the dividing line between covert action (undertaken by the intelligence community) and military operations (undertaken by the Department of Defense) has become increasingly blurred.

Consider, for example, the large amount of intelligence information about adversary systems that is needed to conduct cyberattacks against them. In a targeting context, military collection of the information needed for a cyberattack is essentially indistinguishable from traditional intelligence collection. At the same time, a covert operation undertaken by the intelligence community to influence events in another country may well look like a military operation. Even intelligence collection and exploitation operations may entail some attack activity (and hence appear military-like) in order to gain or preserve access.

Collection activities—presumably including activities requiring cyberattack in some form for their successful execution—would not constitute covert action. Both tapping an adversary's underwater cable to obtain military traffic flows and planting a Trojan horse key logger in an adversary computer system in its ministry of defense would constitute intelligence collection activities, even if such activities were very sensitive.

On the other hand, activities that are intended to influence the conduct, behavior, or actions of an adversary without the involvement of the United States becoming known are covert actions requiring findings if they are not traditional intelligence activities or otherwise exempt, and the dividing line between activities that should be regarded as covert action and those that should not becomes unclear. For example:

- Intelligence preparation of the battlefield is a traditional military activity and thus does not constitute covert action. But a cyberattack may be designed to alter the functionality of an adversary's tactical command and control systems long in advance of actual hostilities on the ground, and thus may be regarded as a covert action.
- Strategic deception conducted under the U.S. military chain of command is a traditional military activity and thus does not constitute covert action. (An example of strategic deception is the attempt to persuade an adversary that an attack will occur in one place when it will actually occur in another.) But a cyberattack may be developed that alters the data streams on which an adversary's intelligence and surveillance capabilities rely, and thus may be regarded as a covert action.
- Collecting telemetry on experimental missile launches is a traditional intelligence collection activity. But a cyberattack may be designed to corrupt or alter the telemetry received by the adversary receiving stations

so that the adversary must redo the test or, even worse, inadvertently use bad data in its R&D efforts, and thus may be regarded as a covert action.

From an administrative or organizational standpoint, command structures may blur the lines between Title 10 authorities (governing the armed forces) and Title 50 authorities (governing the intelligence community). For example, as noted in Chapter 3, the U.S. Strategic Command has responsibility for network warfare—and the Joint Functional Component Command for Network Warfare is commanded by the director of the National Security Agency, an element of the intelligence community. Such blurring requires those in the command structure to be careful about the roles they are playing when they take any given action.

Perhaps the most important point about the distinction between covert action and traditional military activities is that the distinction is essentially irrelevant outside a domestic context. Nations that are the target or subject of an act that they regard as hostile are not likely to care whether the United States classifies it as a covert action or as a military activity. Thus, the entire discussion above relates only to decisions within the U.S. government about how it should organize itself to conduct various activities.

7.3.2 Title III and the Foreign Intelligence Surveillance Act

Domestic electronic surveillance conducted in the United States for purposes of criminal investigation related to any of a list of specifically enumerated offenses is regulated under the federal Wiretap Act of 1968 as amended (also known as “Title III”). Under Title III, law enforcement authorities may seek court authorization to conduct real-time surveillance of electronic communications for these purposes. The court authorization must be issued by a judge who concludes that there is probable cause to believe that a crime relating to one of these enumerated offenses has been, is being, or is about to be committed.

Originally enacted in 1978, the Foreign Intelligence Surveillance Act (FISA) established a framework for the use of “electronic surveillance” conducted to obtain “foreign intelligence information” (defined as information about a foreign power or foreign territory that relates to the national defense, the security, or the conduct of the foreign affairs of the

United States).⁴⁰ For any such surveillance, the statute requires the attorney general and related law enforcement authorities to seek and secure a warrant from a special court known as the Foreign Intelligence Surveillance Court (FISC). A FISC order must specify (among other things) a statement of the means by which the surveillance will be conducted and an indication of the period of time for which the electronic surveillance must be maintained.

Since 1978, FISA has been repeatedly amended to account for new technologies and new concerns about terrorism and civil liberties. The most recent amendments came in 2008. The new statute allows the attorney general and the director of national intelligence to jointly authorize the “targeting of persons reasonably believed to be located outside the United States to acquire foreign intelligence information.” The statute requires the government to adopt “targeting procedures” to meet this goal and “minimization procedures” to avoid the retention or distribution of information concerning U.S. citizens that is obtained from such surveillance. The statute imposes no probable cause requirement for such surveillance, but more restrictive provisions apply when the person targeted overseas is a U.S. national.

Certain cyberexploitations may be regarded as forms of electronic surveillance, and if conducted against U.S. persons or in the United States may under some circumstances be subject to FISA or Title III regulation. Such a cyberexploitation might, for example, require the implantation of software payloads to exfiltrate information surreptitiously. Such information may include important documents relevant for exploitation or information such as login names and passwords that might be useful for conducting a later cyberattack.

It is difficult to speculate on how FISA might be relevant to cyberattacks. But there is at least one documented case of a court-approved Title III warrant being used to authorize a cyberexploitation.⁴¹ On June 12, 2007, an FBI agent filed an affidavit to a magistrate judge in support of an application for court authorization to send a message to a computer used to administer a specific MySpace.com user account. The message was designed to cause this computer to transmit back to the FBI technical data identifying the computer and/or the users of the computer. Whether

⁴⁰ More detailed descriptions of FISA and its impact on intelligence gathering can be found in Elizabeth Bazan, *The Foreign Intelligence Surveillance Act: An Overview of Selected Issues*, Congressional Research Service, Washington D.C., July 7, 2008 (available at www.fas.org/sgp/crs/intel/RL34279.pdf); Elizabeth B. Bazan (ed.), *The Foreign Intelligence Surveillance Act: Overview and Modifications*, Nova Science Publishers, Hauppauge, N.Y., 2008; and Whitfield Diffie and Susan Landau, *Privacy on the Line: The Politics of Wiretapping and Encryption*, Updated and Expanded Edition, MIT Press, Cambridge, Mass., 2007.

⁴¹ See <http://politechbot.com/docs/fbi.cipav.sanders.affidavit.071607.pdf>.

and how often the FISC has approved the use of cyberexploitation, or the nature of such exploitation (if any), is not known from information that is publicly available.

7.3.3 Posse Comitatus

The Posse Comitatus Act (codified at 18 USC 1385), along with administrative action and other related law, prohibits the U.S. armed forces from executing domestic law, unless such actions are explicitly authorized by statute or the U.S. Constitution. (For example, Title 10, Sections 371-381 of the U.S. Code explicitly allow the Department of Defense to provide federal, state, and local police with information (including surveillance and reconnaissance), equipment, and training/expertise. Other legislation has allowed the DOD to assist in matters related to counterterrorism, weapons of mass destruction, and drug trafficking.) Questions arise most often in the context of assistance to civilian police.

Under the Posse Comitatus Act, the Department of Defense would appear to be forbidden from conducting either cyberattack or cyberexploitation in support of domestic law enforcement to enforce domestic law in any context where there was no specific statutory exemption, but would have the authority to conduct such operations domestically if they were part of the exercise of presidential authority to act as commander-in-chief under Article II.

7.3.4 The Computer Fraud and Abuse Act and Other Federal Law

A variety of federal laws, including 18 USC 1030 (the Computer Fraud and Abuse Act, described in Section 5.2) and 18 USC 1029 (dealing with fraud and related activity in connection with access devices), prohibit individuals and corporations from undertaking cyberattack activities. Neither of the statutes mentioned above exempts military agencies from their prohibitions, although the legislative history of each does not suggest that Congress intended it to apply to military operations abroad.

However, the Computer Fraud and Abuse Act may be relevant to possible military cyberattack activities because the various technologies of cyberattack often involve the compromise of third-party computers in order to conceal and otherwise support attack activities against an adversary computer system or network. A party launching a cyberattack—such as the United States—may wish to conceal its identity in such an action. Or, it may wish to augment the computing resources available to it for such purposes at little additional cost.

The issue of public appropriation of private resources depends on whether those private resources are owned by individuals or corporations

in the United States. The law in this area is voluminous and mixed, and the current status of the law about the government's rights to use private computers of Americans without owner permission in the conduct of a cyberattack is quite unclear.

A different analysis, although still murky, applies to the use of private resources owned by individuals or corporations outside the United States. Subsection (f) of 18 USC 1030 (the Computer Fraud and Abuse Act) explicitly states that Section 1030 "does not prohibit any lawfully authorized investigative, protective, or intelligence activity of a law enforcement agency of the United States, a State, or a political subdivision of a State, or of an intelligence agency of the United States." In this context, an activity might be "lawfully authorized" explicitly (as through a warrant granted by the FISC) or implicitly authorized by being undertaken under the legal authority of the President, the bounds of which are evolving and thus not precisely known.

On the presumption that there is no other relevant legislative authority, there appears to be no domestic legislative impediment for the U.S. government to commandeer the computers of private citizens abroad to create a cyberattack capacity for use by the government, perhaps for use in a botnet or perhaps in any attempt to conduct a cyberattack with plausible deniability. Whether such commandeering is legitimate under the international laws of armed conflict is not clear, although the fact that the "zombification" of a computer can leave the computer almost entirely intact and whole for the user's purposes is surely relevant to a LOAC analysis. (As always, whether such actions would be wise or appropriate on policy grounds is an entirely separate matter—this paragraph speaks only to the legal aspect of the issue.)

If none of these approaches worked to allow the U.S. government to assemble a network of computers for a powerful and hard-to-trace cyberattack, there would be the theoretical option to obtain the needed access to large numbers of third-party computers by "renting" them from a private source. But botnets for hire are, as a practical matter, available only from criminals, since it is a criminal act to assemble a botnet in the first place. And although it is not without precedent,⁴² cooperating with or paying criminals to conduct operations relevant to national security is highly problematic, is politically controversial, and may itself be illegal.

Given the leverage available with using third-party computers for cyberattack, government may wish to find other avenues for clarifying the legal landscape for doing so. One approach would be for the U.S.

⁴² One such example of U.S. government cooperation with criminals was the CIA use of Mafia assistance in the attempt to assassinate Fidel Castro in 1960. See "Trying to Kill Fidel Castro," *Washington Post*, June 27, 2007, p. A06.

government to simply ask owners of personal computers for permission to use their computers, or to pay a fee to owners willing to make their computers available for such use.⁴³

Such approaches would obviously eliminate the clandestine nature of such use, but it might well place at the disposal of the U.S. government resources far in excess of what it would otherwise have available. In any event, the committee recognizes that such approaches would be controversial, and it is not advocating them in any way.

7.3.5 The War Powers Resolution

The War Powers Resolution of 1973 was intended to be an assertion of congressional authority relevant to warmaking. A more detailed discussion of the War Powers Resolution is contained in Section 6.2.1.

7.3.6 Executive Order 12333 (United States Intelligence Activities)

Initially promulgated on December 4, 1981, and amended a number of times since then (most recently in July 2008), Executive Order 12333 regulates the conduct of U.S. intelligence activities.⁴⁴ Section 2.2 of Executive Order 12333 sets forth “certain general principles that, in addition to and consistent with applicable laws, are intended to achieve the proper balance between the acquisition of essential information and protection of individual interests.” Using a definition of “United States person” specified in Section 3.4(i) of this order (a United States person is “a United States citizen, an alien known by the intelligence agency concerned to be a permanent resident alien, an unincorporated association substantially composed of United States citizens or permanent resident aliens, or a corporation incorporated in the United States, except for a corporation directed and controlled by a foreign government or governments”), Section 2.3 of Executive Order 12333 establishes constraints on procedures for agencies within the intelligence community to collect, retain or disseminate information concerning United States persons. Section 2.5 requires the attorney general to find probable cause to believe that the U.S. person who is the target of the surveillance is an agent of a foreign power.

⁴³ A partial precedent for using civilian assets for military purposes can be found in the Civil Reserve Air Fleet (CRAF). Under the CRAF program, civilian airlines commit to making available some of their aircraft for military airlift purposes when DOD military aircraft are inadequate to meet a given demand. In return, the government makes peacetime airlift business available to these civilian airlines. See U.S. Air Force Fact Sheet, *Civil Reserve Air Fleet*, available at <http://www.af.mil/factsheets/factsheet.asp?id=173>.

⁴⁴ The full text of Executive Order 12333 as of July 2008 is available at <http://www.tscm.com/EO12333.html.whitehouse.gov/infocus/nationalsecurity/amended12333.pdf>.

U.S. law (including FISA, Title III, state wiretap law, the Electronic Communications Privacy Act, and Executive Order 12333) may restrict the ability of government agencies to collect information within the United States on cyberattacks, just as it places such restrictions on collection on other subjects, including collection of stored information found on the networks of victims, perpetrators, or “hop” sites, as well as collection through wiretapping of communications. The significance of this point is that when a system or network in the United States is the target of a cyberattack, and the perpetrator of that attack is unknown to U.S. authorities (as is almost always the case), collection of that information must be done in accordance with the appropriate and necessary legal authorities.

Absent the consent of the network owners to government collection of the information described above, the legal authorities for law enforcement and (in certain circumstances) counterintelligence provide the broadest basis for such collection. Thus, responsibility for collecting the information required for attack assessment and attribution will normally rest with the FBI (which uniquely possesses both federal law enforcement and counterintelligence collection authorities (including FISA)) and other domestic law enforcement agencies. (Analysis of that information can be—and under the National Infrastructure Protection Center prior to the establishment of the Department of Homeland Security, was—performed jointly by law enforcement, the intelligence community, and military personnel (and by private sector parties if necessary).) Such information is necessary to characterize the nature of an incoming cyberattack, and is of course necessary if any kind of counter-counterattack is to be launched.

In addition, Executive Order 12333 regulates the conduct of covert action by stipulating that “no agency except the CIA (or the Armed Forces of the United States in time of war declared by Congress or during any period covered by a report from the President to the Congress under the War Powers Resolution (87 Stat. 855)1) may conduct any special activity unless the President determines that another agency is more likely to achieve a particular objective,” where “special activities” are defined as “activities conducted in support of national foreign policy objectives abroad which are planned and executed so that the role of the United States Government is not apparent or acknowledged publicly, and functions in support of such activities, but which are not intended to influence United States political processes, public opinion, policies, or media and do not include diplomatic activities or the collection and production of intelligence or related support functions.”

7.4 FOREIGN DOMESTIC LAW

Foreign nations are governed by their own domestic laws governing destructive (that is, attack) computer actions. U.S. cyberattack activities that terminate or transit foreign nations may be subject to such law, though enforcement of those laws may be as a practical matter difficult. Foreign domestic law also has an impact on the ability of the United States to trace the origin of cyberattacks or cyberexploitations directed against the United States—for example, if a certain cyber action is not criminalized in Zendia, Zendian law enforcement agencies may not have the legal authority to investigate it, even if the action is relevant to a cyberattack action against the United States routed by Ruritania through Zendia.

8

Insights from Related Areas

This chapter seeks to contrast and compare cyberconflict with conflict/warfare involving certain other kinds of weapons: nuclear, space, biological, and non-lethal.

8.1 NUCLEAR WEAPONS AND NUCLEAR WAR

As noted in Chapter 6, nuclear history and policy are useful points of departure—framing notions and metaphorical checklists—for understanding issues related to cyberattack, in large part because of the effort that has been devoted to the subject of nuclear conflict over the years. In particular, many questions asked regarding nuclear conflict are relevant to cyberattack, even though the answers to these questions will be very different in the two cases.

Consider first some important differences. Perhaps the most important difference is that the use of a nuclear weapon provides a very important threshold—there is no sense in which the use of even a single nuclear weapon could be regarded as unimportant or trivial. Indeed, a nuclear explosion anywhere in the world, especially one that does damage, is unambiguously detectable even if it is not attributable. By contrast, cyberattacks are being used all the time, not necessarily with government sponsorship or approval, but by criminals and hackers and on a large scale as well. Cyberexploitation also occurs on a large scale, often with no one noticing.

A second key difference relates to attribution. For much of the Cold

War, the bipolar nature of the world—the United States and Soviet Union—would have made it relatively easy for the United States to attribute a nuclear attack. Although a number of other nations had achieved nuclear capabilities as well, these nations either were allies of the United States (and thus could be presumed to not have hostile intent that might lead to the use of nuclear weapons against it) or were generally incapable of striking the United States.

To the extent that the latter proposition is not true, then the United States would have two techniques to determine the identity of an attacking state. First, a network of satellites keeps track of missile launches around the world, and thus the national origin of missile launches can be ascertained. (Missiles launched from the sea are more difficult to attribute.) In addition, radiological analysis of a nuclear explosion's residues might identify the nation responsible for manufacturing the weapon, provided there is on file a record of the radiological "signatures" that would be provided by nuclear weapons from various nations. And nuclear weapons are generally presumed to be under the tight control of the nation's national command authority, and thus the use of a Zendian nuclear weapon could be presumed to be a willful act of the Zendian government.

None of these conditions applies to attribution of cyberattack, as noted in Chapter 2. When it comes to cyberconflict, the world is distinctly *not* bipolar, and indeed nation-states are not the only relevant actors. The true geographic origin of a cyberattack is very difficult to identify. There are no characteristic technical signatures of a given cyberattack that can be unambiguously associated with a specific nation. Finally, a cyberattack cannot be presumed to have been undertaken at the direction of a national government, regardless of where it originates.

Yet another important difference is that the acquisition of nuclear weapons requires an enormous and expensive infrastructure for development, testing, and deployment of those weapons, and thus the threshold for obtaining nuclear weapons is much higher than that for cyberweapons. The elements of such an infrastructure are much easier to observe and identify than the infrastructure needed to acquire cyberweapons. Cyberweapons can be acquired on a small budget behind closed doors using technology that is widely and easily available. In theory, both nuclear weapons and cyberweapons can be purchased, but the sale of a nuclear weapon would be much more visible to national intelligence agencies than the sale of a cyberweapon (some of which can be downloaded for free on the Internet).

Consequently, deterrence through the threat of retaliation has much less credibility for cyberwarfare than for nuclear warfare, a point that in itself is an important difference between cyber and nuclear warfare. (Of course, it is also true that as some of the features of a bipolar adversarial

regime become less relevant or applicable to the state of nuclear affairs today, traditional theories of nuclear deterrence also begin to fray around the edges.¹)

Finally, from an analytical point of view, theories and simulations of escalation dynamics and control have been developed to help understand how a nuclear conflict might unfold—how conflict might transition from non-nuclear to nuclear, the scale and scope of first nuclear use, how such use might lead to subsequent nuclear use, and how nuclear conflict might be terminated. There are few similar theories (at least not in the public literature) about how cyberconflict might unfold, but given the lack of real-world experience with cyberconflict, such theoretical development might well be worthwhile.² Chapter 9 provides a few sketchy speculations on this matter.

There are also a number of similarities between the two domains. From a technical standpoint, one similarity between nuclear weapons and cyberweapons is the superiority of the offense over defense. In both instances, attack operations—i.e., operations that result in destruction or damage—are much easier to undertake than defensive operations, i.e., operations to prevent an attacker from inflicting damage. But the consequences of this similarity are very different in the two cases. In the nuclear domain, this undeniable technical reality has forced the nuclear-armed nations of the world to rely on a strategy of deterrence by threat of retaliation. In the cyber domain, the difficulties of attack attribution leave a comparable threat with far less credibility.

From an operational perspective, military planners have considered the use of nuclear weapons for both strategic and tactical purposes (though debates rage about the wisdom of using nuclear weapons for tactical purposes). In targeting, they can be aimed at adversary military capabilities (counterforce targeting) and societal infrastructure (counter-value targeting). Both can be used in first-use and second-use scenarios. It is technically possible to create automated responses to nuclear attack or cyberattack. At the same time, there are many difficulties in developing a highly reliable and automated assessment regarding both the actual fact of an attack and the appropriate party against which to respond, and thus, the wisdom of such responses in both cases is subject to some considerable question. Finally, both nuclear attack and cyberattack can lead to unintended and unforeseen consequences as well as cascading effects

¹ See, for example, David E. Sanger and Thom Shanker, "U.S. Debates Deterrence for Nuclear Terrorism," *New York Times*, May 8, 2007.

² Of course, the validity of theories of nuclear escalation and control—or of U.S. nuclear doctrine for that matter—has not been tested empirically. Some might regard the net outcome—many untested theories of nuclear conflict and a scarcity of theories of cyberconflict—as more of a similarity between the two domains than a difference.

and liabilities, and attack scenarios involving nuclear weapons and cyber-weapons are highly complex.

From an organizational point of view, both nuclear attack and cyber-attack are complex subjects. They both require deep understanding of technology and policy available only in specialized communities. A great deal of intelligence-based preplanning is needed to construct plausible and realistic attacks with both kinds of weapon, and options can be created in each case for a range of desired effects. Institutionally, both are managed under the U.S. Strategic Command, and the reach of both nuclear weapons and cyberweapons is potentially global.

Other adversary nations and subnational groups are drawn to nuclear weapons and cyberweapons (as well as to other weapons of mass destruction) at least in part because they may serve as equalizers that afford the ability to compete directly but asymmetrically with the United States in conflict situations.

Finally, cyberwarfare and nuclear conflict may be intimately related under some circumstances. For example, the command and control networks used to control nuclear weapons might be targets of cyberattack. A large-scale use of cyberattack weapons that threatens the survival of the targeted nuclear-armed nation could result in its use of nuclear weapons. As noted in Section 6.1.1, U.S. declaratory policy regarding nuclear weapons suggests that the United States could respond to certain kinds of cyberattacks against it with nuclear weapons.

The last point also raises the possibility that the United States might, under some circumstances, choose to refrain from using cyberattacks that are intended to have large-scale, society-damaging effects, at least against nuclear-armed states. This point is explored further in Section 9.2 on escalation dynamics and control.

8.2 SPACE

Operations in space provide a few lessons for understanding cyberattack and cyberexploitation. (For purposes of this discussion, operations in space are limited to operations involving satellites.)

Satellites can be attacked in a number of ways. They can be destroyed by kinetic impact (such as by a direct-ascent missile) or by directed energy weapons (either land-based or space-based) that cause the satellite to overheat or that destroy on-board optical or infrared sensors. Such “hard-kill” options render a satellite permanently inoperative.

“Soft-kill” options interfere with the satellite’s operation, rendering it non-functional, but in a reversible manner. One might, for example, jam its command uplink so that it cannot receive commands from the ground. In the absence of such commands, a satellite might not be able to execute

a given mission or it might even drift out of position. A satellite may use an unencrypted command link, so that an adversary could manipulate the satellite's functions. A more fanciful approach for soft kill might entail the unfurling of a large aluminized Mylar bag around the adversary satellite that prevented commands from reaching it or from using its on-board sensors. Attacks on the ground control stations of a satellite could also render a satellite non-functional, although a nation that relied on satellites heavily would be likely to have backup ground stations for such contingencies.

Apart from attacks on ground stations, attacks on satellites would almost certainly be non-lethal—there would be no military value in attacking a crewed space vehicle. But an attack on an important satellite would undoubtedly have strategic impact. That is, if undertaken before kinetic conflict had broken out, such an attack would be regarded by the satellite-owning nation as a major provocation, and it undoubtedly would qualify as a hostile “use of force” against that nation. If it were undertaken after kinetic conflict had broken out, it would inevitably be regarded as a significant escalation of the conflict.

Some kinds of cyberattack share some of these characteristics. As noted in Chapter 2, the immediate effects of cyberattack are almost always non-lethal, but the consequences of certain kinds of cyberattack, such as attacks on the infrastructure of a nation, could have large-scale strategic impact. And, depending on how they were configured, cyberattacks may result in hard kill or soft kill of their targets.

Intelligence collection is another point of legal similarity between operations in space and cyber operations. Today, there is broad international acceptance of the principle that reconnaissance satellites can transit freely and without prior approval over national boundaries. Similarly, cyberexploitations have not traditionally been regarded as violations of international law.³

8.3 BIOLOGICAL WEAPONS

Biological weapons and cyberweapons share a number of similarities—indeed, the term “virus” as an instrument of cyberattack was adopted in recognition of a mode of large-scale attack with certain similarities to how biological viruses spread and attack hosts.

It is helpful to consider biological weapons and cyberweapons with respect to two categories—characteristics of the weapons themselves,

³ Public opinion and perceptions of these two acts are quite different—there is little public outcry against the reconnaissance satellites of other nations directed against the United States, but there is a great deal of public outcry against cyberexploitations directed against the United States.

and characteristics of the infrastructure needed to produce and use such weapons.

One major similarity of biological weapons and cyberweapons is that the release of the weapons agent and/or its effects may well not be immediately detectable. A biological virus can be released quietly in a crowded football stadium (no loud explosions), and people will become sick days later. A computer virus can be released on the Internet without notice, and can lie dormant on targeted computers for extended periods without anyone noticing symptoms such as degradation in computer performance and so on. And its effects will be noticed only if the virus is triggered.

In both cases, the weapon can replicate without requiring human intervention—biological viruses or bacteria can multiply; computer viruses and worms copy themselves. One result is that weapons effects may continue after and beyond the point of the initial attack. The disease caused by a bioweapon may propagate through secondary contagion (i.e., human carriers of a disease), whereas the effects of a cyberattack may propagate or cascade beyond the point of the initial attack (as other computers are attacked).

It is possible for cyberattack weapons to be selective about the targets on which they inflict damage—for example, a virus or a worm may be configured to cause damage only to selected systems even if it propagates to a large number of systems. In principle, biological weapons might be tailored to cause disease only in individuals with a certain biological signature, even if it infects others without causing disease.⁴

Furthermore, much of society is constructed in ways that enhance the efficacy of biological weapons and cyberweapons. The effectiveness of biological weapons is enhanced by high population densities in urban areas and by poor health care and public health/epidemiological reporting systems; the effectiveness of cyberweapons is enhanced by high dependence on interconnected information technology and a lack of concerted attention to cybersecurity on a societal scale.

“Blowback” from biological weapons and from cyberweapons is an important concern. Blowback refers to the phenomenon in which a weapon loosed on an enemy blows back against the weapons user. A

⁴ See, for example, British Medical Association, *Biotechnology, Weapons and Humanity*, Harwood Academic Publishers, Amsterdam, the Netherlands, 1999; and Claire M. Fraser and Malcolm R. Dando, “Genomics and Future Biological Weapons: The Need for Preventive Action by the Biomedical Community,” *Nature Genetics* 29(3):253-256, November 2001, available at <http://cmbi.bjmu.edu.cn/news/report/2001/insight-anthrax/feature/Genomics%20and%20future%20biological%20weapons.pdf>. The issue of such targeted weapons was raised as early as 1970 in the professional military literature. See Carl Larson, “Ethnic Weapons,” *Military Review* 50(11):3-11, November 1970, available at <http://usacac.army.mil/CAC/Repository/Materials/MilitaryReview-197011300001-DOC.pdf>.

biological virus used by Zendia against Ruritania may, in an unknown period of time, affect Zendian citizens en masse. Similarly, a Zendian computer virus targeted against Ruritanian computers may eventually infect Zendian computers.

8.4 NON-LETHAL WEAPONS

Non-lethal weapons constitute yet another area from which some relevant insights may be gleaned. Box 8.1 provides some illustrative examples of non-lethal weapons.

A preliminary similarity is the struggle over appropriate terminology regarding non-lethal weapons, a struggle that reprises the analogous issue

BOX 8.1 Non-lethal Weapons—Illustrative Examples

Traditional Instruments

- Night sticks and truncheons
- Water cannons that shoot jets of water at high pressure
- Rubber bullets
- Tear gas
- Pepper spray
- Dogs

Today's Instruments

- Tasers
- Flashbangs (which create loud sounds or sudden bursts of light or bad smells)
- Projectile netting
- Carbon filaments (for use against electrical grids, to short out switching stations)
- Loud music (e.g., Noriega and the use of Nancy Sinatra's "These Boots Are Made for Walking")

Future Systems

- Sticky or slippery foams
- Non-nuclear electromagnetic pulse weapons for use against vehicles
- Malodorants
- Sound cannons (for projecting loud sounds at standoff distances, e.g., against small boats)
- Active denial systems (e.g., a vehicle-mounted millimeter-wave heat ray that creates intense heat pain through clothing without actually causing burns)

raised in Chapter 1 about information warfare, information operations, cyber operations, and so on. Non-lethal weapons have come to designate a category of weapons that are explicitly designed and primarily employed so as to incapacitate personnel or materiel while minimizing fatalities, permanent injury to personnel, and undesired damage to property and the environment. But there are no assurances or guarantees of non-lethality—no matter how carefully designed or carefully used, a given “non-lethal” weapon may result in fatalities if it is used against a particularly vulnerable person. One proposed alternative calls such weapons “less lethal,” but objections have been raised to that term as well as indicating that such weapons would be used to create undead zombies. A clumsy term might be “weapons with significantly reduced probability of lethality,” but clumsy terms are hard to use in discourse.

One policy issue raised by non-lethal weapons involves a seductive quality about them that has the potential of lulling users into a sense of complacency about their use. For example, the *New York Times* reported on a study by the sheriff’s office in Orange County, Florida, in which the officers on patrol were all equipped with tasers and were trained to use them.⁵ One immediate effect was that the number of citizen fatalities due to police action decreased dramatically—the hoped-for effect. A second immediate effect was a dramatic increase in the frequency of police use of force overall. That is, prior to the introduction of tasers, the police might not have used force in any way—they might have talked the person down or waited him out or might have found some way to resolve the matter without using force. But with tasers in hand, they were more willing to use force (that is, to use a weapon) than before. This effect had not been anticipated.

A similar issue arises with cyberweapons, which are also non-lethal with respect to their immediate effects. Perhaps more importantly, they offer the opportunity to avoid the use of traditional lethal weapons—and for policy makers seeking to take actions short of the use of such weapons, they may be similarly seductive. That is, if policy makers see them as weapons without lethal effects, they may be more inclined to favor options calling for their use⁶ or to specify rules of engagement for using them in the field that are more permissive than would be the case for kinetic weapons.

⁵ Alex Berenson, “As Police Use of Tasers Soars, Questions Over Safety Emerge,” *New York Times*, July 18, 2004.

⁶ The search for actions that are “short of force” is apparent in almost every instance in which economic sanctions are proposed against some nation. That is, economic sanctions are almost always the first actively adversarial action taken against nations that offend the international order.

A related point is whether the existence of non-lethal weapons (or perhaps cyberweapons) places legal or moral/ethical obligations to use them before lethal weapons are used. Similar questions have arisen in the context of using smart versus dumb bombs. It can be argued that both morality and the law of armed conflict requires the use of the weapons that are the most discriminating in their ability to minimize collateral damage—by this argument, a military force would be required to use smart bombs (that is, weapons that can be more accurately aimed) before it used dumb bombs (weapons that are less discriminate in their destruction). To date, the United States and other nations have resisted any such argument, but these issues may recur from time to time in the future as weapons become even more discriminate.

Finally, both law enforcement agencies and the Department of Defense have equities and interests in the area of non-lethal weapons. But their interests and priorities are different, and it is hard to point to a single authoritative voice within the U.S government on the subject. Similarly, the U.S. Air Force and the National Security Agency (and perhaps other intelligence agencies as well) also have an interest in cyberattack and offensive cyber operations, and the different interests and priorities of these institutions will have to be reconciled.

Speculations on the Dynamics of Cyberconflict

9.1 DETERRENCE AND CYBERCONFLICT

To what extent is deterrence of cyberconflict possible? How might a nation's cyberweapons be useful in deterring an adversary's cyberattack?

In the language of defense policy, deterrence is an often-used and highly elastic concept, and it is hard to find an authoritative statement of its precise meaning. For purposes of this document, the definition provided by the U.S. Strategic Command is a reasonable starting point:¹

Deterrence [seeks to] convince adversaries not to take actions that threaten U.S. vital interests by means of decisive influence over their decision-making. Decisive influence is achieved by credibly threatening to deny benefits and/or impose costs, while encouraging restraint by convincing the actor that restraint will result in an acceptable outcome.

The threat "to impose costs" is the foundation of classical deterrence, more specifically deterrence by threat of retaliation or punishment. This concept was the underpinning of U.S. nuclear policy toward the Soviet Union during the Cold War, and continues to be central to the reality of dealing with other nuclear states today. At the same time, an opponent that can be deterred by the threat of imposing costs is, almost by defini-

¹ Deterrence Operations: Joint Operating Concept, Version 2.0, December 2006, available at http://www.dtic.mil/futurejointwarfare/concepts/do_joc_v20.doc.

tion, a rational opponent—that is, one who can calculate that the costs of a certain action outweigh the gains possible from taking that action and thus does not take that action. But it is well known and widely understood that some actors are not rational in this sense of the term. Such “non-rational” actors may indeed be able to make rational cost-benefit calculations, and still take the “non-rational” course of action because of political or religious ideology, a belief in luck, or even insanity.

The threat “to deny benefits” is the rationale for the deployment of defensive capabilities—capabilities that can interfere with the success of an attack. Antiballistic missile defenses, for example, are intended to prevent hostile ballistic missiles from striking friendly targets. Chemical protective suits are intended to reduce the effectiveness of chemical weapons against friendly forces. Offensive counter air operations are intended to destroy hostile aircraft before those aircraft take off to conduct attacks against friendly targets and territory.

A refinement on the concept of deterrence as described above is the notion of tailored deterrence—deterrence “tailored” to specific adversaries in specific strategic contexts. For example, the U.S. Strategic Command notes that:

Exercising decisive influence over the decision calculations of adversary decision-makers requires an understanding of their unique and distinct identities, values, perceptions, and decision-making processes, and of how these factors are likely to manifest themselves in specific strategic contexts of importance to the US and its allies. Specific state and non-state adversaries thus require deterrence strategies and operations tailored to address their unique decision-making attributes and characteristics under a variety of strategically relevant circumstances. Such tailored deterrence strategies and operations should be developed, planned, and implemented with reference to specific deterrence objectives that identify who we seek to deter from taking what action(s), under what conditions (i.e., Deter adversary X from taking action Y, under Z circumstances).

Box 9.1 describes the key questions of tailored deterrence.

It remains an open question as to whether the concepts of deterrence are relevant when applied to the domain of cyberconflict per se (that is, cyberconflict without reference to conflict in physical domains). For example, a credible threat to impose costs requires knowledge of the party on which the costs should be imposed—and as discussed in Chapter 2, attribution of a cyberattack is a very difficult and time-consuming—and perhaps insoluble—problem.

Moreover, even if the adversary is known, and known to be a specific nation-state, the costs to be imposed must be judged by the adversary as greater than the gain that might result from his aggressive actions. Thus, the United States must be able to identify cyber targets in or of the adver-

BOX 9.1 Tailored Deterrence

Tailoring an approach to deterrence requires answering four questions as described below. Specific answers to all four questions would represent a specific tailoring.

1. *Who is being deterred?* By definition, deterrence is intended to influence an adversary's decision-making process in such a way that the adversary chooses to refrain from taking an action that is undesirable to the United States. Further, the mechanisms through which deterrence can operate depend strongly on the party the United States is trying to influence. Possible answers for the question "Who is being deterred?" include:

- The national leadership of an adversary nation
- Leaders of subnational groups
- Private citizens of the adversary nation

2. *What is the undesirable action to be deterred?* Depending on the undesirable action to be deterred, different threats and different targets (below) might be required. Possible actions to be deterred include:

- Nuclear attack
- Attack with biological or chemical weapons
- Attack with conventional forces
- Cyberattack
- Adversary interventions in other locales

3. *What threat is the basis for the deterrent?* By definition, deterrence involves a threat of some kind. A common approach to determining the deterrent threat is the threat of "in-kind" action—deterring *X* action by an adversary calls for threatening to do *X* to the adversary. But in-kind action is not inherently necessary for deterrence, and much of the U.S. approach to deterrence has explicitly called for threats that are not symmetric. For example, the United States has long reserved the right to use nuclear weapons against an overwhelming conventional attack or against attacks using biological or chemical weapons. Some of the possible threats that might be used to deter an adversary include:

- Nuclear attack
- Attack with biological or chemical weapons
- Conventional attack
- Cyberattack
- Economic or diplomatic pressure

4. *What is the target of the U.S. threat?* A threat must be directed at a target or targets whose loss would be important enough to the adversary decision maker to make him refrain from taking the undesirable action. Some possible targets might include:

- Nuclear forces
- Biological or chemical weapon forces or stockpiles
- Conventional forces
- Leadership
- Key industries
- Economic infrastructure
- Population

sary nation whose loss would be costly to the adversary, and it must be able to attack them with high confidence of success.

In a nation that is not highly dependent on information technology, such assets would be hard to find. Even if the nation did have valuable information technology assets, specific individual targets (perhaps numbering in the dozens or hundreds—a wild guess!) most valuable to the adversary are likely to be very well protected against cyberattack. The civilian IT infrastructure at large may be less well protected, but large-scale attacks on such infrastructure raise ethical and moral questions about targeting civilians. The military IT infrastructure could be targeted as well, but the degree to which it is well protected may be unknown to the attacker (see discussion in Chapter 2 regarding intelligence requirements for successful focused cyberattacks).

In addition, an attacker that launches a cyberattack should also be expected to take action to change its own defensive posture just prior to doing so. As discussed in Chapter 2, much can be done to invalidate an adversary's intelligence preparations, which are necessary for discriminating counterattacks. And since the attacker knows when he will launch the attack, he can create a window during which his defensive posture will be stronger. The window would last only as long as it would take for new intelligence efforts to collect the necessary information, but it would likely be long enough to forestall immediate retaliation.

A threat to deny benefits to a cyberattacker also lacks credibility in certain important ways. In principle, defensive technologies to harden targets against cyberattacks can be deployed, raising the difficulty of attacking them. But decades of experience suggest that deploying these technologies and making effective use of them on a society-wide basis to improve the overall cybersecurity posture of a nation is difficult indeed. And there is virtually no prospect of being able to reduce a cyberattacker's capabilities through offensive action, because of the ease with which cyberattack weapons can be acquired. Thus, counterforce capabilities—which in the nuclear domain have been justified in large part as necessary to reduce the threat posed by an adversary's nuclear weapons—do not exist in any meaningful way in contemplating cyberconflict.²

How do the considerations above change if, as in the real world, the states involved also have kinetic capabilities, which may include nuclear weapons, and physical vulnerabilities? That is, each side could, in principle, use kinetic weapons to attack physical targets, and these targets might be military or dual purpose in nature as long as they are legitimate targets

² This statement is NOT intended to indicate acceptance or rejection of the counterforce argument in the nuclear domain—it is only to say that regardless of whether the counterforce argument is valid in the nuclear domain, it has little validity in the cyber domain.

under LOAC. Because a transition from cyber-only conflict to kinetic conflict would likely constitute an escalation (and would in any case make the conflict more overt), this point is discussed in more detail below.

9.2 ESCALATORY DYNAMICS OF CYBERCONFLICT BETWEEN NATION-STATES

The escalatory dynamics of conflict model how a conflict, once started, might evolve. Of interest are issues such as what activities or events might set a cyberconflict into motion, what the responses to those activities or events might be, how each side might observe and understand those responses, whether responses would necessarily be “in kind,” how different kinds of state might respond differently, and so on. What follows below are some speculations on some of the factors that might influence the evolution of a cyberconflict.

The actors involved are presumed to be nation-states with significant kinetic and cyber capabilities at their disposal, and the situation in question is one of open tension and high rhetoric between two states that have traditionally been rivals. Important questions to be addressed (summarized in Box 9.2) are discussed in the remainder of this section, but the discussion is intended to raise issues rather than to answer questions.

9.2.1 Crisis Stability

Where kinetic weapons are concerned, crisis stability refers to that condition in which neither side has incentives to attack first. Crisis stability is especially important for nuclear weapons, where the existence of an invulnerable submarine-based nuclear missile force means that an adversary could not escape retaliation no matter how devastating or successful a first strike it could launch against the United States. Where cyberweapons are concerned, there is no conceivable way for a nation to eliminate or even significantly degrade the cyberattack capability of another nation.³ But the question remains whether a second-strike cyberattack capability is the enabling condition for crisis stability in cyberspace.

A related question is that of incentives for preemption. Suppose that preemptive attacks by Ruritania on Zendia are undertaken in order to prevent (or at least to blunt) an impending attack by Zendia on Ruritania.

³ Even in the case of a nuclear electromagnetic pulse attack directed against the electronic equipment in another nation (Zendia), there is no reason to assume that all of Zendia’s cyberattack capabilities are necessarily resident within Zendia’s boundaries. Because cyberattacks can originate from anywhere, some of Zendia’s cyberattack capabilities might have been deployed in other nations—indeed, some Zendian attack agents might already have been clandestinely deployed in U.S. systems.

BOX 9.2 Questions About the Escalatory Dynamics of Cyberconflict Between Nation-States

Crisis Stability

- What is the analog of crisis stability in cyberconflict?
- What are the incentives for preemptive cyberattack?

Escalation Control and Management

- How can intentions be signaled to an adversary in conflict?
- How can cyberconflict between nations be limited to conflict in cyberspace?
- How should cyberattack be scoped and targeted so that it does not lead an adversary to escalate a conflict into kinetic conflict?
- How can a modestly scoped cyberattack conducted by a government be differentiated from the background cyberattacks that are going on all of the time?
- How can the scale and scope of a commensurate response be ascertained?

Complications Introduced by Patriotic Hackers

- How can “free-lance” activities on the part of patriotic hackers be handled?

Incentives for Self-restraint in Escalation

- What are the incentives for self-restraint in escalating cyberconflict?

Termination of Cyberconflict

- What does it mean to terminate a cyberconflict?

If Zendia is planning a cyberattack on Ruritania, a preemptive cyberattack on Zendia cannot do much to destroy Zendia’s attack capability; at best, Ruritania’s preemptive attack on Zendia might tie up Zendia’s personnel skilled in cyber operations. On the other hand, it is hard to imagine circumstances in which Ruritania would realize that Zendia was planning an attack, as preparations for launching a cyberattack are likely to be invisible for the most part.

A second relevant scenario is one in which Zendia is planning a kinetic attack on Ruritania. Intelligence information, such as photographs of troop movements, might well indicate that preparations for such an attack were being made. And under these circumstances, Ruritania might well choose to launch a preemptive cyberattack against Zendia, with the intent of delaying and disrupting Zendia’s preparations for its own (that is, Zendia’s) kinetic attack.

9.2.2 Escalation Control and Management

In a time of tension or crisis, national leaders are often understandably concerned about inadvertent escalation. For example, Nation *A* does *X*, expecting Nation *B* to do *Y* in response. But in fact, Nation *B* unexpectedly does *Z*, where *Z* is a much more escalatory action than *Y*. Or Nation *A* may do *X*, expecting it to be seen as a minor action intended only to show mild displeasure and thinking that Nation *B* will do *Y* in response, where *Y* is also a relatively mild action. But due to a variety of circumstances, Nation *B* sees *X* as a major escalatory action and responds accordingly with *Z*, an action that is much more significant than *Y*. Nation *A* perceives *Z* as being way out of proportion, and in turn escalates accordingly.

9.2.2.1 Signaling Intentions Through Cyberconflict

Nothing in the alphabet of options above is specific to cyberconflict—such issues have been an important part of crisis management for a long time. But managing such issues may well be more difficult for cyberconflict than for other kinds of conflict. One reason is the constant background of cyberattack activity. Reports arrive hourly and daily of cyberattacks of one kind or another on U.S. computer systems and networks, and the vast majority of these attacks do not have the significance of a serious cyberattack launched by a party determined to do harm to the United States. Indeed, the intent underlying a given cyberattack may not have a military or a strategic character at all. Organized crime may launch a cyberattack for profit-making purposes. A teenage hacking club may launch a cyberattack out of curiosity or for vandalism purposes.

A dearth of historical experience with nations or terrorists using cyberattack against the United States further complicates efforts at understanding what an adversary might hope to gain by launching a cyberattack. And other nations are in a similar position, lacking the experience and facing the same background of cyberattacks. In the absence of contact with cyberattackers (and sometimes even in the presence of such contact), determining intent is likely to be difficult, and may rest heavily on inferences made on the basis of whatever attack attribution is possible.

Thus, attempts to send signals to an adversary through limited and constrained military actions—problematic even in kinetic warfare—are likely to be even more problematic when cyberattacks are involved.

9.2.2.2 Preventing Cyberconflict from Transitioning to Physical Space

If national command authorities decide to retaliate in response to a cyberattack, an important question is whether retaliation must be based

on a “tit-for-tat” response. Assuming the perpetrator of a cyberattack is known to be a hostile nation, there is no reason in principle that the retaliation to a hostile cyberattack could not be a kinetic attack against the interests of that hostile nation—that is, allowing a kinetic response to a cyberattack expands the range of options available to the victim. An extreme case is that in the event of a cyberattack of sufficient scale and duration to threaten a nation’s ability to function as a modern society, the attacked nation might choose to respond with kinetic force to the nation causing such problems. On the other hand, the attacked nation may have an interest in refraining from a kinetic response—for example, it may believe that a kinetic response would be too provocative and might result in an undesired escalation of the conflict.

Decision makers may also see cyberattacks as instruments to be used in the early stages of a conflict (cf. Section 3.2). National decision makers considering a cyberattack (whether in response or as a first use) appear to have incentives to refrain from conducting cyberattacks that might induce a strong kinetic reaction unless kinetic conflict had already broken out. The obvious approach would be to conduct cyberattacks that are in some sense smaller, modest in result, targeted selectively against less provocative targets, and perhaps more reversible. (The similarity of such an approach to escalation control in other kinds of conflict is not accidental, and it has all of the corresponding complexities and the uncertainties.)

There is no reason to suppose that hackers and criminal elements will moderate their activities in times of crisis or conflict (see also Section 9.2.3 regarding patriotic hackers). Thus, if a cyberattack is intended to send a signal from the United States to Zendia, how is Zendia to recognize that signal? Overtly taking credit for such an attack goes only so far, especially given uncertain communications in times of tension or war, and the near certainty of less-than-responsible behavior on the part of one or both sides.

Finally, it seems likely that escalation issues would play out differently if the other nation(s) involved are near-peer competitors or not. Escalation to physical conflict is of less of concern to the United States if the nation has weak conventional forces and/or is a non-nuclear state. But a nation with nuclear weapons, or even with strong conventional forces in a position to damage U.S. allies, is another matter entirely, and relationships with such states may well need to be specially managed, paying particular attention to how escalation may be viewed, managed, and controlled, and most importantly, how miscalculation, misperception, or outright error may affect an adversary’s response.

9.2.2.3 Determining the Impact and Magnitude of Cyber Response

If an adversary conducts a cyberattack against the United States, a first question for U.S. decision makers will be knowledge of the attack's impact and magnitude. Such knowledge is necessary to inform an appropriate U.S. response. (If, for example, the United States wishes to make a commensurate response, it needs to know what parameters of the incoming attack would characterize a commensurate response.)

But in many kinds of cyberattack, the magnitude of the impact of the first cyberattack will be uncertain at first, and may remain so for a considerable period of time. Decision makers may then be caught between two challenges—a policy need to respond quickly and the technical fact that it may be necessary to wait until more information about impact and damage can be obtained. (As noted in Section 2.5, these tensions are especially challenging in the context of active defense.)

Decision makers often feel intense pressure to “do something” immediately after the onset of a crisis, and sometimes such pressure is warranted by the facts and circumstances of the situation. On the other hand, the lack of immediate information may prompt decision makers to take a worst-case view of the attack and thus to assume that the worst that might have happened was indeed what actually happened. Such a situation has obvious potential for inappropriate and unintended escalation.

9.2.3 Complications Introduced by Patriotic Hackers

Past experience strongly indicates that conflict or increased tension between two nations will result in the “patriotic hackers” of both nations (and perhaps their allies) taking action intended to harass or damage the other side. Such activities are not under the direct control of the national government, and as discussed in Section 7.2.3.3 may well interfere with the efforts of that government to manage the crisis vis-à-vis the other side.⁴ Indeed, the government of a targeted nation is likely to believe that a cyberattack conducted on it is the result of deliberate adversarial action rather than the actions of “unauthorized” parties. Thus, unauthorized activities of the patriotic hackers of Zendia against the United States may lead the United States to believe that the Zendian government has launched a cyberattack against it. A U.S. cyberattack against Zendia may be seen by the Zendian government as a cyber first strike against it.

Yet another complication involving patriotic hackers is the possibility that they might be directed by, inspired by, or tolerated by their govern-

⁴ Such activities also have some potential for complicating the operational efforts of that government—for example, because cyberattacks against the same target may interfere with each other.

ment (or a rogue section within it), but in ways in which the government's hand is not easily visible. Under such circumstances, hostile acts with damaging consequences could continue to occur (with corresponding benefits to the nation responsible) despite official denials. At the very least, the possibility that patriotic hackers may be operating could act as a plausible cover for government-sponsored cyberattacks, even if there were in fact no patriotic hackers doing anything.

9.2.4 Incentives for Self-restraint in Escalation

One set of incentives is based on concerns about an adversary's response to escalation. Understanding this set of incentives is necessarily based on a sense of what kinds of offensive cyber actions, whether cyberattack or cyberexploitation that might be mistaken for cyberattack, might lead to what kinds of adversary responses (in cyberspace or in physical space). In this regard, an essential difference between cyberattack and the use of nuclear, chemical, biological, or space weapons is readily apparent—the initial use of any nuclear, chemical, biological or space weapon, regardless of how it is used, would constitute an escalation of a conflict under almost any circumstances. By contrast, whether a given cyberattack (or conventional kinetic attack for that matter) would be regarded as an escalation depends on the nature of the operation—the nature of the target(s), their geographic locations, their strategic significance, and so on.

A second set of incentives is based on concerns about “blowback”—the possibility that a cyberattack launched by the United States against Zendian computers might somehow affect U.S. computers at a later time. Understanding the likelihood of blowback will require a complex mix of technical insight and intelligence information.

9.2.5 Termination of Cyberconflict

How could the United States indicate to Zendia that it was no longer engaging in cyberattacks against it? Given that a cyberattack might well involve the placement of hardware and/or software agents within the Zendian IT infrastructure (both civilian and military), would the United States direct such agents to self-destruct? Would it inform Zendia of the IT penetrations and compromises it had made? On what basis would the Zendian government believe a claim by the United States that it had issued such a directive? (And, of course, all of the same questions apply in reverse as well.)

On the other hand, such actions may be more analogous to cleanup and recovery efforts after a kinetic war. Conflict termination in a kinetic

war means that both sides stop shooting at each other—and refrain from taking further destructive actions. This point suggests that software and/or hardware agents within an adversary’s IT infrastructure must be designed so that they are under the positive control of the launching nation—and thus that fully autonomous agents are inconsistent with positive control. In addition, an attacker may need to keep careful track of where these agents are implanted, so that subsequent “cyber de-mining” operations are possible when hostilities have terminated.

9.2.6 The Role of Transparency

Where kinetic weapons are concerned, transparency and confidence-building measures such as adherence to mutually agreed “rules of the road” for naval ships at sea, prenotification of large troop movements, and non-interference with national technical means of verification have been used to promote stability and mutual understanding about a potential adversary’s intent.

Secrecy surrounding cyberattack policy works against transparency. In addition, military operations on land, sea, and air are easily distinguishable from most non-military movements, whereas it is likely to be difficult to distinguish between military and non-military cyber operations.

9.2.7 Catalytic Cyberconflict

Catalytic conflict refers to the phenomenon in which a third party instigates conflict between two other parties. These parties could be nation-states or subnational groups, such as terrorist groups. The canonical scenario is one in which the instigator attacks either Zendia or Ruritania in such a way that Zendia attributes the attack to Ruritania, or vice versa. To increase confidence in the success of initiating a catalytic war, the instigator might attack both parties, seeking to fool each party into thinking that the other party was responsible.

As also noted in Section 2.4.2, high-confidence attribution of a cyberattack under all circumstances is arguably very problematic, and an instigator would find it by comparison very easy to deceive each party about the attacker’s identity. Thus, a catalytic attack could be very plausibly executed. In addition, if a state of tension already exists between the United States and Zendia, both U.S. and Zendian leaders will be predisposed toward thinking the worst about each other—and thus may be less likely to exercise due diligence in carefully attributing a cyberattack. A Ruritanian might thus choose just such a time to conduct a catalytic cyberattack.

9.3 CYBERCONFLICT BETWEEN THE UNITED STATES AND NON-STATE ACTORS

Competition with nation-states is not the only kind of conflict that might involve the United States. For example, the United States might be the target of a cyberattack by a non-state party (such as a terrorist group).

A terrorist group, by definition, does not operate as a nation-state, and there would inevitably be difficulties in identifying the relevant terrorist group (many terrorist groups would surely like to be able to conduct a cyberattack against the United States), thus complicating the “impose costs” strategy.

In addition, if the terrorist group were operating under the auspices of a failed state, a cyber counterattack would be likely to find few suitable targets in the failed state and thus would have little impact. (Kinetic counterattack might be feasible, as indicated by the experience of the United States in attacking Afghanistan immediately after the September 11, 2001, terrorist attacks on the World Trade Center and the Pentagon.) If the terrorist group were operating under the unwitting cover of another state, all of the attribution problems described in Section 2.4.2 would apply, and the discussion in Section 7.2.1.2 on the merits of self-defense in neutral territory would be relevant.

Criminal groups conducting cyberattacks for illicit monetary gain are also important non-state actors. If they operate across national boundaries, law enforcement efforts to shut down the operations of such groups are likely to take a long time, if they are successful at all. Thus, one might plausibly consider, in addition to the usual law enforcement efforts, a different response paradigm that could call for cyberattacks to terminate or attenuate their activities.

Against non-state parties, deterrence by retaliation may be particularly ineffective. First, a non-state group may be particularly difficult to identify. A lack of identification means uncertainty about the appropriate focus of any retaliatory action, and also that the decision-making calculus of the non-state group is likely to be poorly understood. Second, a non-state group is likely to have few if any information technology assets that can be targeted. Third, some groups (such as organized hacker groups) regard counterattacks as a challenge to be welcomed rather than something to be feared. A criminal group might react very strongly to a counterattack by a much stronger cyberattack than was initially launched. Fourth, a non-state group such as a terrorist or insurgent group might

seek to *provoke* cyber retaliation in order to galvanize public support for it or to antagonize the public against the United States.⁵

A particularly challenging problem is the prospect of an extended cyber-guerilla campaign against the United States by a non-state actor (perhaps with state sponsorship) operating globally over periods of month or years. In many ways, coping with such a campaign is similar to the physical space kinetic analog of an extended terrorist campaign against the United States. For example:

- The set of possible targets is nearly infinite, suggesting that hardening every possible target against attack is an implausible strategy.
- Knowing the adversary's value calculus (what the adversary values and how he values it) is fraught with uncertainty, which makes strategies based on deterrence by retaliation less effective as a policy tool. For example, if the United States cannot determine assets that the adversary values, credible retaliation is impossible to threaten. If the adversary-valued assets are in a friendly state, attacking those assets might have negative repercussions.
- Penetration of the entities responsible for hostile actions against the United States (that is, turning an insider) is likely to be very problematic because of the difficulty of identifying an insider with whom to engage.
- A continuing campaign, whether kinetic or cyber, could be very effective in instilling fear, terror, and uncertainty in the population regardless of the actual level of damage being inflicted.⁶

There may also be factors that differentiate the situation as compared to the kinetic analog. For example, kinetic terrorism usually has dramatically visible effects, whereas the effects of a sustained guerilla campaign of cyberattacks may be far less visible, especially against the background noise of daily cyberattacks from myriad sources. Uncertainty and an undermining of confidence in information technology may be the most likely result of a cyber campaign. Also, state sponsorship may pro-

⁵ The notion of provoking U.S. retaliation as a technique for gaining the sympathies of the Islamic world at large is a basic tenet of Al Qaeda's strategy against the United States. See, for example, Rohan Gunaratna, *Inside Al Qaeda's Global Network of Terror*, Columbia University Press, New York, 2002.

⁶ For example, in the D.C. sniper case of 2002, 10 people were shot over a period of 3 weeks (Jessica Reaves, "People of the Week: John Muhammad and John Malvo," *Time*, October 24, 2002, available at <http://www.time.com/time/nation/article/0,8599,384284,00.html>). Inhabitants of the greater D.C. area were terrorized and fearful because of the sniper, despite the fact that there were 18 "traditional" homicides during that time (Susan Kim, "Fear Lingers in DC Area," *Disaster News Network*, November 12, 2002, available at <http://www.disasternews.net/news/article.php?articleid=57>).

vide those responsible with access to intelligence that could amplify the potency of guerilla cyberattacks.

How might the United States combat such a campaign? Although it is probably possible to harden genuinely critical targets of cyberattack and thereby make a truly devastating attack more difficult, the number of possible lucrative targets is large. Thus, such a campaign could still be expected to have some non-trivial degree of success. Because the locus of an attack can be shifted arbitrarily and essentially instantaneously, active threat neutralization would provide at best transient relief, if any at all. Moreover, enlisting the assistance of foreign national authorities is problematic because a shifting geographic locus can easily negate the effectiveness of any assistance offered.

An alternative to the methods described above is to use techniques such as deception and infiltration coupled with covert cyberattack. Deception might be used to induce operatives to violate operational cybersecurity by opening themselves to cyberattack (e.g., to visit putatively useful websites that might be able to infect visitors with a selectively acting Trojan horse). Infiltration may be difficult (as described above) but would have to be a priority effort. But whether this alternative would in fact be more effective against a cyberterrorist group is an open question.

9.4 THE POLITICAL SIDE OF ESCALATION

The discussions in previous sections in this chapter address escalation dynamics primarily from a military standpoint. Yet escalation dynamics inevitably has a political and psychological component that must not be overlooked.

For example, Section 2.5 (on active defense) points out that U.S. cyberattacks undertaken under the rubric of active defense may not be perceived by others as innocent acts of self-defense, even if they are intended by the United States as such. While in most conflicts, both sides claim that they are acting in self-defense, cyberconflicts are a particularly messy domain in which to air and judge such claims. Another possible misperception may arise from intelligence collection activities that might involve cyberattack techniques. As noted in Section 2.6.1, the tools needed to conduct a cyberexploitation may not be very different from those needed to conduct a cyberattack. On the other hand, a nation's tolerance for being the target of a cyberattack may be much lower than its tolerance for being the target of a cyberexploitation.

Thus, consider the political ramifications in the following troublesome scenarios:

- Zendia might believe that it has been attacked deliberately by the United States even when the United States has not done so. Indeed, because of the ongoing nature of various attack-like activities (e.g., hacking and other activities) against the computer systems and networks of most nations, the Zendian conclusion that Zendian computer systems are being attacked is certainly true. Attribution of such an attack is a different matter, and because hard evidence for attribution is difficult to obtain, the Zendian government might make inferences about the likelihood of U.S. involvement by giving more weight to a general understanding of U.S. policy and posture toward it than might be warranted by the specific facts and circumstances of the situation. Evidence that appears to confirm U.S. involvement will be easy to find, whether or not the United States is actually involved, and the lack of U.S.-specific “fingerprints” can easily be attributed to U.S. technological superiority in conducting such attacks.

- An active defense undertaken by the United States of its systems and networks against Zendia could have significant political consequences. For example, even if the United States had technical evidence that was incontrovertible (and it never is) pointing to the Zendian government, the Zendians could still deny that they had launched such an attack—and in the court of world opinion, the Zendian denial could carry some weight when considered against past U.S. assertions regarding similar issues. That is, U.S. cyberattacks (counter-cyberattacks, to be precise) undertaken under the rubric of active defense may not be perceived as innocent acts of self-defense, even if they are. The result could be a flurry of charges and countercharges that would further muddy the waters and escalate the level of political tension and mistrust.

- The United States plants a software agent in a Zendian military system but does not activate it (cf. Section 2.2.4). Zendia (being attacked) may well regard the hostile action as beginning at the moment the U.S. agent is planted, whereas the United States may believe that the hostile action begins only when the agent is activated.

- The United States launches a cyberattack against a Zendian military factory, but the direct damage from this attack is not visible to the naked eye (Section 2.3.1.1). Without CNN images of smoking holes in the ground or troops on the move, an outside observer must weigh competing claims without tangible evidence one way or the other. Under such circumstances, the reputations of the different parties in the eyes of each other are likely to play a much larger political role.

- The United States plants software agents in some of Zendia’s critical networks to collect intelligence information. These agents are designed to be reprogrammable in place—that is, the United States can update these agents with new capabilities. During a time of crisis, Zendian authorities discover some of these agents and learn that they have been present for a

while, that they are sending back to the United States very sensitive information, and that their capabilities can be changed on a moment's notice. Even if no harmful action has yet been taken, it is entirely possible that Zendia would see itself as being the target of a U.S. cyberattack.

- The United States is the target of a cyberattack against its air traffic control system that results in a number of airplane crashes and several hundred deaths. Initially, no definitive technical attribution can be made regarding the perpetrator of the attack, but in a matter of weeks, an all-source attribution—depending on somewhat uncertain human and signals intelligence—suggests that the perpetrator could be Zendia. The United States decides on a mixed kinetic and cyber response against Zendia but must persuade allies and the rest of the world that its attack on Zendia is in fact justified.

- Tensions between the United States and Zendia are high, even though diplomats are trying to defuse them. Over a relatively short period of time, Zendia conducts a number of cyberexploitations against a variety of computer systems and networks important to the U.S. military. Some of these activities are successful in compromising some sensitive but unclassified information, but the systems and networks in question do not experience any apparent functional degradation. However, in keeping with common press usage, U.S. news reports of these activities indicate that 300 Zendian “cyberattacks” have taken place against the U.S. military. In turn, these reports inflame passions in the United States, leading to significant pressures on the U.S. National Command Authority to respond aggressively against Zendia.

Factors such as the ones described above suggest that factors other than those dictated by military or legal necessity play important roles in escalation dynamics, if nothing else because they can strongly affect the perceptions of decision makers on either side.

Alternative Futures

As described in Chapters 3-5, the stance of the United States toward cyberattack against adversary foreign nations is one that puts no constraints on its use apart from those imposed by the law of armed conflict and related customary international law. But such a stance is not the only possible one, and from time to time proposals emerge that, if adopted, would constrain activities related to cyberattack for some or all nations, including the United States. This chapter explores some of the issues that arise in considering such proposals, but does not take a stand one way or another on their inherent desirability.

10.1 REGULATORY REGIMES—BASIC PRINCIPLES

The laws of armed conflict acknowledge an inevitability to conflict and seek to put restraints on what might otherwise be unrestrained behavior. In addition, nations that may engage in armed hostilities with one another sometimes enter into legal regimes that regulate the development, testing, production, acquisition, deployment, or use of certain kinds of weapons. Such regimes—generically arms control regimes—are generally regarded as having some mix of three broad purposes: to reduce the likelihood that conflict will occur, to reduce the destructiveness of any conflict that does occur, and to reduce the costs associated with the acquisition of the weapons that are the subject of the agreement or with defense against those weapons.

Arms control agreements can be bilateral between two nations (such

as the Strategic Arms Reduction Treaties between the United States and the Soviet Union/Russia) or multilateral among multiple nations (such as the Limited Test Ban Treaty signed and ratified by 94 nations). They can be cast formally as treaties, informally as memorandums of understanding, or even more informally as coordinated unilateral policies. They may place limits on the acquisition of certain kinds of weapons, where acquisition can be understood to mean research, development, testing, production, or some combination thereof (e.g., a ban on the development, testing, production, and deployment of intermediate-range ballistic missiles); on the deployment of certain weapons (e.g., no nuclear weapons in space); on the use of such weapons (e.g., prohibitions on the use of laser weapons specifically designed, as their sole combat function or as one of their combat functions, to cause permanent blindness to unenhanced vision¹); or on the circumstances of weapons use (e.g., an agreement to refrain from “first use” of nuclear weapons).

In many cases, and especially when they involve the use of certain kinds of weapons, arms control agreements are seen by the signatories as confidence-building measures, that is, actions taken or not taken that are intended to provide a potential adversary with reassurances that some other action is not hostile in intent. For example:

- The United States and the Soviet Union maintained a “hot line” to facilitate direct contact between the respective national leaders during times of crisis on the theory that direct contact would be valuable in reducing misunderstanding about national activities that were ongoing or imminent.
- The United States and the Soviet Union signed an agreement in 1989 that bound each side to take steps to prevent interference with command and control networks in a manner that could cause harm to personnel or damage to equipment of the armed forces of the other side.²
- The United States and Russia have another agreement to notify each other 24 hours in advance prior to the launch of a strategic ballistic missile.³ The intent of this agreement is to reassure the other party that

¹ Note that the United States has not ratified the Protocol on Blinding Laser Weapons (Protocol IV to the Convention on Certain Conventional Weapons).

² Agreement of the Government of the United States of America and the Government of the Union of Soviet Socialist Republics on the Prevention of Dangerous Military Activities, June 1989, available at http://en.wikisource.org/wiki/Prevention_of_Dangerous_Military_Activities_Agreement.

³ Agreement Between the United States of America and the Union of Soviet Socialist Republics on Notifications of Launches of Intercontinental Ballistic Missiles and Submarine-Launched Ballistic Missiles, available at <http://www.state.gov/t/ac/trt/4714.htm>.

in the event that a strategic ballistic missile is launched by the first party, such a launch is not misunderstood as a prelude to hostilities.

- The United States and Russia have agreed to various measures to reduce the likelihood of an incident at sea between the naval forces of the two countries, and to reduce the likelihood of escalation in the event that one occurred. Such measures include steps to avoid ship collisions, avoiding maneuvers in areas of heavy sea traffic, requiring surveillance ships to maintain a safe distance from the object of investigation, refraining from simulating attacks at the other party's ships, and so on.⁴

Arms control agreements often contain measures to enhance verification—a process by which one signatory can develop confidence that the other side is indeed living up to its obligations under the agreement. Some agreements, such as confidence-building measures, are self-verifying—each nation undertakes to enact or engage in those measures when they are called for in the agreement, and if the nation does not do so when appropriate, the other nation draws whatever conclusions it may draw about the other side's intentions. Other agreements provide for the use of “national technical means” (i.e., various technical intelligence assets) and/or various kinds of inspections to verify compliance. Still other agreements make no provision for verification at all (such as the Biological Weapons Convention), but nevertheless serve as statements regarding international norms of acceptable conduct that constrain, at the very least, the declaratory policies of the signatories to be consistent with the agreements in question.

Many critics of arms control agreements point to a lack of verification provisions as a fatal flaw in an agreement. They argue that when the United States is party to such an agreement, it is invariably bound by both the spirit and the letter of the agreement, but that the other party—usually an adversary or a potential adversary of the United States—is likely to violate the agreement in the absence of adequate verification provisions, thus leaving the United States at a relative disadvantage.

These basic principles of arms control regimes can be applied to understanding possible approaches to developing international agreements regulating cyberattack.

⁴ Agreement Between the Government of the United States of America and the Government of the Union of Soviet Socialist Republics on the Prevention of Incidents On and Over the High Seas, available at <http://www.state.gov/t/ac/trt/4791.htm>.

10.2 REGULATORY REGIMES FOR CYBERATTACK

10.2.1 Direct Approaches Based on Traditional Arms Control

What purposes could be served by a regulatory regime for cyber-attack? Traditional arms control theory generally indicates that three broad purposes could be served in principle,⁵ presuming that the restrictions of the regime are observed by all signatories:

- *Reducing the likelihood that conflict will occur.* Confidence-building measures—arrangements in which signatory parties agree to refrain from or to notify other signatories prior to conducting certain activities that might be viewed as hostile or escalatory or to communicate directly with each other during times of tension or crisis—are explicitly intended to reduce the likelihood of conflict due to accident or misunderstanding. In addition, agreements to eschew the use of cyberattack may have some value in reducing the likelihood of kinetic conflict in those cases in which cyberattack is a necessary prelude to a kinetic attack.⁶

- *Reducing the destructiveness of any conflict that does occur.* Limitations on targeting cyberattack weapons could prevent damage to the prohibited entities, presuming that the scope of a cyberattack can be delimited with confidence. Moreover, limiting damage to those entities might prevent escalation from occurring—and such escalation could include escalation to kinetic or even nuclear conflict. Reducing destructiveness might also facilitate a more rapid cessation of cyberhostilities.

- *Reducing financial costs.* Limitations on acquisition of weapons for cyberattack would not have a significant impact on financial costs, simply because these weapons are so inexpensive in the first place. Nor would a particular adversary's agreement to refrain from conducting cyberattack relieve the United States from needing to defend against other nations or subnational entities that could use such weapons.

Given the possibilities for cyberattack to disrupt national economies or to distort the activities of individual companies as well (especially large

⁵ These three purposes can be found in Thomas C. Schelling and Morton H. Halperin, *Strategy and Arms Control*, Pergamon-Brassey's, Washington, D.C., 1985.

⁶ Such cases may well be rare. If and when they exist, they are based on the idea that cyberattack is to be used for shaping battlefield conditions and introducing delay and disruption into adversary planning. (See, for example, the discussion in Chapter 9 (escalation), Chapter 3 (how the United States is likely to use information warfare should it become necessary), and Section 10.3 (regarding China).) If a kinetic attack requires the kind of battlefield conditions and the delay or disruption that only a cyberattack can provide, then the kinetic attack might be inhibited. On the other hand, an adversary may well have alternative (non-cyber) means for accomplishing these tasks.

companies that are very important to a nation), a regulatory regime for cyberattack might also reduce the likelihood of economic warfare using this military tool.

As an example of an international agreement involving the use of cyberattack, Davis Brown, a former deputy staff judge advocate of the U.S. Defense Information Systems Agency,⁷ has proposed to extend the law of armed conflict to account explicitly for the use of information systems in armed conflict (Box 10.1).

The argument for the United States entering into an international agreement regarding cyberattack⁸ is based on the notion that the United States would be relatively worse off than any other nation if all nations could carry out cyberattacks without restriction because the United States is significantly more dependent on information technology than any other nation that is likely to be involved in a major cyberconflict. Whether this relative disadvantage will endure over the long term depends on whether the dependence of other nations on information technology is increasing more rapidly than that of the United States—but it is undeniable from any perspective that the United States would have much to lose in all-out cyberconflict whether or not that loss would be greater or less than that suffered by an adversary.

In this view, an agreement regarding cyberattack weapons is based in large part on a desire to delegitimize such use against the United States, precisely because the United States has so much to lose from a large-scale cyberconflict. Conversely, aggressive pursuit of cyberattack capabilities by the United States is seen as legitimizing cyberattack as a military weapon and indeed as encouraging other nations to develop such capabilities for use against the United States and its interests. Others argue that other nations need no prodding from the United States to develop cyberattack weapons for use against it, and that adversary development of such weapons is inevitable regardless of what the United States chooses to do in this arena.

Another benefit of a formal agreement regarding use of cyberattack is that it can help to make explicit many of the concerns that military operators will have (or, at least, should have) in using cyberattack as an operational weapon. If certain operational practices are prohibited, questions about whether or not an operator can engage in those practice are easier to resolve.

⁷ Davis Brown, "A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict," *Harvard International Law Journal* 47(1):179-221, Winter 2006.

⁸ A possible platform on which such an agreement might be constructed is the Convention on Conventional Weapons, an international agreement that regulates a number of individual weapons, including lasers intended to blind humans.

BOX 10.1 An Illustrative Draft Convention Regulating the Use of Information Systems in Armed Conflict

Davis Brown, a former deputy staff judge advocate of the U.S. Defense Information Systems Agency, proposes to extend the law of armed conflict to account explicitly for the use of information systems in armed conflict. Brown accepts conventional LOAC as the initial point of departure and is guided by the principle that an act that violates LOAC if carried out by conventional means also violates LOAC if carried out by cyberattack.

Under Brown's proposal:

- The activities of patriotic hackers against an adversary would be prohibited.
- The military forces conducting cyberattack should not be commingled with civilians in their workplaces.
 - Cyberattacks on dual-use infrastructure (e.g., railroads, communications centers, pipelines) are legitimate as long as the military advantage gained by attacking such targets outweighs the harm to civilians.
 - The use of cyberattack to attack civilian infrastructure or targets whose destruction would cause severe environmental damage would be prohibited.
 - The use of cyberattack weapons whose impact is indiscriminate—that cannot distinguish between military and civilian targets—or that cannot self-destruct or be rendered harmless after hostilities terminate would be prohibited.
 - Cyberattacks on the military payroll system or on non-combatant families of military personnel or posting the Social Security numbers of individual servicemen and servicewomen to increase their vulnerability to identity theft would be prohibited.
 - Active threat neutralization would be permitted even if it involved damage to innocent third parties whose computers had been compromised, if passive defense was insufficient to defend against the threat.
 - Only certain kinds of false identities would be prohibited. These prohibited false identities would include masquerading as an official in the government or armed forces of the target state or of any third state, and masquerading as originating from any third state, or as originating with any medical or religious establishment in any location.
 - Belligerents would be forbidden to use for military purposes domain names or computer systems associated with neutral nations, to launch cyberattacks from computer systems in neutral states, or to take control of neutral systems in order to conduct cyberattacks.

SOURCE: Adapted from Davis Brown, "A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict," *Harvard International Law Journal* 47(1):179-221, Winter 2006.

What about the verifiability of any such agreement? Consider first the feasibility of verifying an agreement to refrain from acquiring cyber-attack capabilities. Many factors suggest that such an agreement would not be verifiable in any meaningful way. The technology—hardware technology—of certain kinds of cyberattack is easily available at Staples, Best Buy, and Dell.com, and its acquisition cannot be limited. The knowledge needed to conduct such cyberattacks is more difficult to acquire but is also available on the Internet, to say nothing of knowledge developed by sophisticated computer scientists. Code—software tools—to carry out cyberattack can be transmitted over the Internet and reproduced trivially, and is available from many sources. Restricting the development of the expertise needed to conduct cyberattacks is equally implausible, because the expertise needed to develop defenses against cyberattacks is intimately related to the expertise needed to develop cyberattacks themselves. Nor would any acceptable inspection regime have a meaningful chance to find software-based cyberattack weapons. Finally, the human and technical infrastructure needed to conduct cyberattack would be much smaller than (and could easily be embedded within) that needed to conduct cyberdefense on a large scale, and thus could be easily hidden.

An agreement might also involve restrictions on the use of cyber-attack weapons. For example, signatories might agree to refrain from striking at national financial systems or power grids, much as nations might avoid targeting hospitals in a kinetic attack, or to refrain from using lasers intended to blind soldiers. In order to facilitate the non-attack of such facilities, nations might agree to take measures to electronically identify systems as being associated with prohibited targets,⁹ much as the “robots.txt” protocol today is used to signal search engines to refrain from indexing a given website.¹⁰ A more limited agreement might obligate signatories to refrain from first-use cyberattacks on national financial systems or power grids.

Obviously, an attacker can ignore such electronic indicators, just as a kinetic attacker can ignore red crosses painted on the sides of ambulances in times of war. Moreover, such agreements are not “verifiable” in advance, in the sense that no amount of information collected before a conflict can guarantee that restrictions on use will be observed during conflict. But such agreements do create international norms regarding the acceptability of such behavior, and they do something to inhibit training

⁹ Davis Brown, “A Proposal for an International Convention to Regulate the Use of Information Systems in Armed Conflict,” *Harvard International Law Journal* 47(1):179-221, Winter 2006.

¹⁰ For more on this protocol, see <http://www.robotstxt.org/>.

that calls for such use. The threat of reciprocal use during conflict may also serve as a deterrent to first use.

In the case of cyberattack, restrictions on use are complicated by many factors. For example, subnational groups under the nominal jurisdiction of a signatory may take actions independently of the government. A nation's military forces may refrain from targeting the power grids of an adversary, but patriotic hackers or terrorist groups on that nation's soil might do so without explicit government approval. Thus, compliance with such an agreement might entail a somewhat bizarre scenario in which two nations are in conflict, perhaps kinetic conflict, but each is simultaneously conducting actions (perhaps involving law enforcement) to suppress subnational cyber actions intended to advance their respective causes. On the other hand, such agreements are likely to be more effective prior to the onset of conflict, because a signatory would have incentives to take suppressing actions in order to avoid undue and unwanted escalation.

Moreover, arms control agreements have in the past presumed a state monopoly on the arms being regulated. But in the case of tools that might be used for cyberattack, the private sector owns and operates much of the infrastructure through which cyberattacks might be conducted. Indeed, the behavior of individual citizens might be directly affected by a traditional arms control agreement—and the degree of intrusiveness on the behavior of individuals and the private sector more generally might be large indeed depending on the nature of the agreement.

Furthermore, the technology with which to conduct cyberattacks is most assuredly not exclusively or even mostly controlled by governments. Private citizens (hackers) conduct many cyberattacks on their own every day. Non-state actors such as terrorist groups or transnational criminal organizations could develop significant cyberattack capabilities as well, but would be unlikely to adhere to any agreement between the United States and any of the nations that might harbor them.¹¹ Under such circumstances, domestic laws in the relevant nations may be the only legal means of regulating the activities of such parties (and even then, the effectiveness of domestic laws depends on the availability of some enforcement mechanism, which may not be present in some of these nations).

Another complication is the functional similarity between cyberexploitation against an adversary's information systems and cyberattack against those information systems. A cyberexploitation may well be interpreted by the target as a damaging or destructive act (or at least the prelude to such an act), and yet to eschew the former actions would be to

¹¹ Jason Barkham, "Information Warfare and International Law on the Use of Force," *New York University International Law and Politics* 34:57-113, 2001.

contradict what amounts to standard operating procedure for essentially all nations.

A final complication discussed in this report arises from the difficulty of tracing cyberattacks to their ultimate origin. If the ultimate origin of a cyberattack can be concealed successfully, holding the violator of an agreement accountable becomes problematic. One technological approach is to deploy a supporting infrastructure more capable than that of today which could support a “use control” regime—a more technologically secure network on a different physical infrastructure whose use would be restricted to those willing to subject themselves to a more constrained regime regarding behavior (e.g., who would agree to be strongly authenticated) and classified as critical to national well-being. But deploying such an infrastructure has many potential drawbacks, such as preventing any connection, physical or logical, to the regular Internet through which cyberattacks might be launched; retaining the economies of the present-day Internet; and preventing the compromise of the strongly authenticated machines.

Agreements might also take place among allies, though in such instances they may take the form of what might be called coordinated unilateral declaratory policies. For example, the NATO nations could collectively agree to refrain from using large-scale cyberattacks against the entire critical infrastructure of an adversary nation as a matter of declaratory policy. Any such agreement—or more precisely, discussions leading to such an agreement—will inevitably stimulate dialogue and debate regarding the topic of cyberattack.

Finally, the history of arms control agreements is that they often suffer from the too-early/too-late problem. That is, the desirability of an agreement may be anticipated, but the technology, doctrine, and so on are not well developed at the time, so it is premature to enter into an agreement. Then technology and doctrine advance rapidly, and before it is widely realized, it has become too late to enter into an agreement because the potential signatories to such an agreement have so much at stake in using the weapons that would be controlled by the putative agreement.

10.2.2 Indirect Approaches Based on Regulation of Non-military Domains

The United States has been a party to many international agreements that are not arms control agreements. For example, nations have sometimes agreed on the need to protect some area of international activity such as airline transport, telecommunications, maritime activities, and so on, and also on standards for such protection. They may declare certain purposes collectively with regard to a given area of activity on which

they agree, often in the form of a multilateral treaty, and then establish consensus-based multilateral institutions (generally referred to as “specialized agencies” composed of experts rather than politicians) to which to delegate (subject to continuous review) the task of implementing those agreed purposes.

Sofaer and Goodman argue that it has been easier to obtain agreement among the nations involved on standards and methods for regulating the civilian (commercial) aspects of a given activity than to obtain agreement on standards and methods for regulating the military (governmental) aspects of the same activity.¹² For example, civil aviation is regulated internationally through agencies that have promulgated numerous agreements and regulations, all by consensus. Over the years, some precedents, and some forms of regulation, have been established, again largely by consensus, that have enhanced the protection of civilian aviation and reduced the uncertainties regarding governmental (military) aviation. A similar pattern of international regulation has resulted in increased maritime safety.

In both areas, states have agreed to criminalize terrorist attacks, and to prosecute or extradite violators. These commitments have not uniformly been kept, but security has been enhanced in these areas of international commerce because of the virtually universal support given to protecting these activities from identified threats.

Sofaer and Goodman proposed a draft multilateral treaty that would have initiated a similar process to help improve cybersecurity internationally, even though it would have initially excluded any direct application of rules and standards developed to the national security activities of member states. The proposed treaty would have included:

- Agreed principles on the use and protection of cyberspace;
- Maximum emphasis on protecting the system, rather than on preventing its use for socially unacceptable objectives such as pornography;
- Agreement of all parties to cooperate in preventing, prosecuting, and cooperating against improper conduct by any non-government group;
 - Maximum coverage, so as to limit use of “rogue” territories as bases for attacks;
 - A program to develop cyber capacities of developing states; and
 - Substantial involvement and authority given to the private sector in developing and approving standards.

¹² Abraham D. Sofaer and Seymour E. Goodman, *A Proposal for an International Convention on Cyber Crime and Terrorism*, Center for International Security and Cooperation, Stanford University, August 2000.

However, the U.S. government rejected the concept of a multilateral treaty with comprehensive aims in favor of a narrower treaty with European allies limited to establishing certain cyber-system crimes and securing commitments for cooperation in dealing with those activities—the Convention on Cybercrime described in Section 7.2.6.

Sofaer and Goodman argue that the approach they propose would, over the long run, provide greater, broad-based international support for a meaningful international cybersecurity regime than will result from a more limited approach.

10.3 FOREIGN PERSPECTIVES ON CYBERATTACK

The potential impact of cyberattacks on a nation's defense posture has not gone unnoticed in other nations or in the world community. For example, in September 1998, then-Russian foreign minister Igor Ivanov wrote to Kofi Annan, United Nations secretary-general, warning that the effect of information weapons "may be comparable to that of weapons of mass destruction."¹³ Likely in response to that letter, the United Nations General Assembly subsequently considered an item entitled "Developments in the Field of Information and Telecommunications in the Context of International Security"¹⁴ and has adopted a resolution on this topic several times since then. These resolutions have variously called on member states to further promote the multilateral consideration of existing and potential threats in the information security field, as well as possible measures to limit emerging threats, consistent with the need to preserve the free flow of information. In addition, they have invited all member states to inform the secretary-general of their views on several topics, including a "general appreciation of the issues of information security"; "definition of basic notions related to information security that would include unauthorized interference with or misuse of information and telecommunications systems and information resources"; and "relevant international concepts aimed at strengthening the security of global information and telecommunications systems."¹⁵

Although several member states have indeed submitted views on this topic, the efforts of the General Assembly have been spearheaded by the

¹³ See letter from Ivanov to Annan, September 30, 1998, available at http://www.un.org/ga/search/view_doc.asp?symbol=A/C.1/53/3&Lang=E.

¹⁴ UN Document A/RES/53/70, "Developments in the Field of Information and Telecommunications in the Context of International Security," January 4, 1999, available at <http://daccess-ods.un.org/TMP/7333411.html>.

¹⁵ *United Nations Disarmament Handbook*, United Nations Publications, New York City, 2004, available at <http://www.un.org/disarmament/HomePage/ODAPublications/Yearbook/2004/Html/Ch%20V6.html>.

Russian Federation, and it is not coincidental that important source documents contributing to the UN General Assembly discussion of the topic are authored by senior scholars and others from Russia.

For example, some Russian thinkers have noted the potentially strategic significance of information warfare and have connected the consequences of information attacks to potentially nuclear responses:

From a military point of view, the use of information warfare means against Russia or its armed forces will categorically not be considered a non-military phase of a conflict, whether there were casualties or not. . . . Considering the possible catastrophic consequences of the use of strategic information warfare means by an enemy, whether on economic or state command and control systems, or on the combat potential of the armed forces, . . . Russia retains the right to use nuclear weapons first against the means and forces of information warfare, and then against the aggressor state itself.¹⁶

In stating its views on the subject of information security to the United Nations, Russia defined information war as “confrontation between States in the information area for the purpose of damaging information systems, processes and resources and vital structures, undermining political, economic and social systems as well as the massive psychological manipulation of a population in order to destabilize society and the State.” Information weapons were regarded as the “ways and means used for the purpose of damaging the information resources, processes and systems of a State, exerting an adverse influence, through information, on the defence, administrative, political, social, economic and other vital systems of a State, as well as the massive psychological manipulation of a population in order to destabilize society and the State.”

The Russian Federation has set forth to the United Nations a document articulating what it describes as “Principles of International Information Security.”¹⁷ (Selected principles are listed in Box 10.2.) The document appears to be intended as a draft resolution of the United Nations General Assembly. The intent of the Russian statement of principles appears to be an outright prohibition on the national development, creation, and use of tools for cyberattack (Principle II.a), on interfering with or unlawfully

¹⁶ V. I. Tsymbal, “Kontseptsiya `Informatsionnoy Voyny” (Concept of Information Warfare), speech given at a Russian-U.S. conference, “Evolving Post-Cold War National Security Issues,” Moscow, September 12-14, 1995, p. 7, cited in Timothy L. Thomas, “Deterring Information Warfare: A New Strategic Challenge,” *Parameters* 26(Winter):82, 1996-1997.

¹⁷ United Nations General Assembly A/55/140, *Developments in the Field of Information and Telecommunications in the Context of International Security*, Fifty-fifth session, Item 69 of the provisional agenda, July 20, 2000, available at <http://www.un.org/documents/ga/docs/55/a55140.pdf>.

BOX 10.2 Selected Russian Principles of International Information Security

Principle II

States shall strive to restrict threats in the field of international information security and with that end in view shall refrain from:

- (a) The development, creation and use of means of influencing or damaging another State's information resources and systems;
- (b) The deliberate use of information to influence another State's vital structures;
- (c) The use of information to undermine the political, economic and social system of other States, or to engage in the psychological manipulation of a population in order to destabilize society;
- (d) Unauthorized interference in information and telecommunications systems and information resources, as well their unlawful use;
- (e) Actions tending to establish domination or control in the information area;
- (f) Preventing access to the most recent information technologies and the creation of conditions of technological dependency in the information field to the detriment of other States;
- (g) Encouraging the activities of international terrorist, extremist or criminal associations, organizations, groups or individual law breakers that pose a threat to the information resources and vital structures of States;
- (h) Formulating and adopting plans or doctrines envisaging the possibility of waging information wars and capable of instigating an arms race as well as causing tension in relations between States and specifically giving rise to information wars;
- (i) The use of information technologies and tools to the detriment of fundamental human rights and freedoms in the field of information;
- (j) The transboundary dissemination of information in contravention of the principles and norms of international law and of the domestic legislation of specific countries;
- (k) The manipulation of information flows, disinformation and the concealment of information in order to corrupt the psychological and spiritual environment of society, and erode traditional cultural, moral, ethical and aesthetic values;
- (l) Expansion in the field of information and the acquisition of control over the national information and telecommunications infrastructures of another State, including the conditions for their operation in the international information area.

Principle III

The United Nations and appropriate agencies of the United Nations system shall promote international cooperation for the purpose of limiting threats in the field of international information security and creating, for that purpose, an international legal basis to:

- (a) Identify the defining features of information wars and to classify them;
- (b) Identify the characteristic features of information weapons, and of tools that may be regarded as information weapons, and to classify them;
- (c) Restrict traffic in information weapons;
- (d) Prohibit the development, dissemination or use of information weapons;
- (e) Prevent the threat of the outbreak of information wars;
- (f) Recognize the danger of using information weapons against vital structures as being comparable to the threat of use of weapons of mass destruction;
- (g) Create conditions for the equitable and safe international exchange of information based on the generally recognized rules and principles of international law;
- (h) Prevent the use of information technologies and tools for terrorist or other criminal purposes;
- (i) Prevent the use of information technologies and tools to influence social consciousness in order to destabilize society and the State;
- (j) Develop a procedure for the exchange of information on and the prevention of unauthorized transboundary influence through information;
- (k) Create an international monitoring system for tracking threats that may arise in the information field;
- (l) Create a mechanism for monitoring compliance with the conditions of the international information security regime;
- (m) Create a mechanism to resolve conflict situations in the area of information security;
- (n) Create an international system for the certification of information and telecommunications technologies and tools (including software and hardware) with a view to guaranteeing their information security;
- (o) Develop a system of international cooperation among law enforcement agencies with a view to preventing and suppressing crime in the information area;
- (p) Harmonize, on a voluntary basis, national legislation in order to ensure information security.

SOURCE: United Nations General Assembly A/55/140, *Developments in the Field of Information and Telecommunications in the Context of International Security*, Fifty-fifth session, Item 69 of the provisional agenda, July 10, 2000, available at <http://www.un.org/documents/ga/docs/55/a55140.pdf>.

using information systems or resources (II.d), and on developing plans or military doctrines intended to wage “information wars” (II.h). To support these goals, the Russian statement calls on the United Nations to create an international legal basis to identify the characteristic features of information weapons, and to classify them (III.b); to restrict traffic in information weapons (III.c); and to prohibit the development, dissemination, or use of information weapons (III.d).

Official Russian position statements to the United Nations notwithstanding, it is widely believed that Russia is fully engaged in, or at least developing, the capability for launching cyberattacks, regardless of its UN stance.

China’s view on the topic of cyberconflict appears to be radically different from that of the Russian Federation. One analyst of Chinese military forces identifies 10 “information operations” methods that the Chinese anticipate using:¹⁸

- Planting information mines,
- Conducting information reconnaissance,
- Changing network data,
- Releasing information bombs,
- Dumping information garbage,
- Disseminating propaganda,
- Applying information deception,
- Releasing clone information,
- Organizing information defense,
- Establishing network spy stations.

China apparently sees great value in acquiring information warfare capabilities and developing facility in their use, and indeed sees information warfare as an equalizer in potential military conflicts with a technologically superior adversary such as the United States.¹⁹ For example, Mulvenon argues that the Chinese see information warfare against the information systems of the U.S. military as a way to degrade and delay the mobilization of U.S. forces and/or their deployment to Taiwan in the event of a crisis over that territory.²⁰

¹⁸ Timothy L. Thomas, “China’s Electronic Strategies,” *Military Review* (May-June), 2001. Available at http://leav-www.army.mil/fmso/documents/china_electric/china_electric.htm.

¹⁹ James C. Mulvenon, “The PLA and Information Warfare,” in James C. Mulvenon (ed.), *The People’s Liberation Army in the Information Age*, Conference Proceedings, The RAND Corporation, 1998.

²⁰ Two PLA authors explicitly endorse what they call “asymmetric information offensives.” See Wang Jianghuai and Lin Dong, “Viewing Our Army’s Quality Building from the

The Eligible Receiver exercise of 1997 underscores this point. According to *Government Executive*,²¹ the exercise—designed to expose weaknesses in computer security in unclassified DOD computer systems using off-the-shelf technology and software downloaded from hacker websites—demonstrated how hackers might disrupt troop deployments.

But the Chinese also believe that political and economic targets as well as military targets are fair game for information warfare. Indeed, disruption of these institutions is an important element in demoralizing an adversary and reducing its will to fight, and so the Chinese view it as entirely reasonable to attack financial systems, power generation and transmission facilities, and other elements of critical infrastructure as part of conflict with another nation (whether or not that conflict has become kinetic).

Finally, the Chinese also see information warfare as a way of enabling the citizenry to participate in a military conflict,²² in which any citizen with a computer can participate in information warfare against an adversary. Indeed, according to Thomas, the information warfare mission is an ideal one for the reserve military forces of China, which can enlist many individuals who are not qualified or eligible to be frontline soldiers.

The Chinese perspective suggests that the Chinese are likely to view any attempt to restrict the use of cyberattack as a way to undermine one of China's few advantages in competing militarily with an adversary such as the United States.

Perspective of What Information Warfare Demands," *Jiefangjun bao*, March 3, 1998, p. 6, in FBIS-CHI-98-072, March 13, 1998; cited in Mulvenon, 1998.

²¹ Katherine McIntire Peters, "Information Insecurity," *Government Executive*, April 1, 1999, available at <http://www.govexec.com/features/0499/0499s1.htm>.

²² Timothy L. Thomas, "Like Adding Wings to the Tiger: Chinese Information War Theory and Practice," Foreign Military Studies Office, Fort Leavenworth, Kans. Undated publication, available at <http://fms.leavenworth.army.mil/documents/chinaiw.htm>.

Appendixes

Appendix A

Biographies of Committee Members and Staff

COMMITTEE MEMBERS

William A. Owens, *Co-chair*, is the chair and CEO of AEA Holdings based in Hong Kong. He retired as vice chair and chief executive officer of Nortel on November 15, 2005. Before joining Nortel in 2004, Admiral Owens was chief executive officer and chair of Teledesic LLC and president, chief operating officer, and vice chair of Science Applications International Corporation (SAIC). Before joining SAIC, he was vice chairman of the Joint Chiefs of Staff and the second ranking military officer in the United States. He had responsibility for the reorganization and restructuring of the armed forces in the post-Cold War era. Widely recognized for bringing commercial high technology into the Department of Defense for military applications, the admiral was the architect of the Revolution in Military Affairs (RMA), an advanced systems technology approach to military operations that is the most significant change in the system of requirements, budgets, and technology for the four armed forces since World War II. From 1991 to 1993, he was the deputy chief of Naval Operations for Resources, Warfare Requirements and Assessments. Admiral Owens served as commander of the U.S. Sixth Fleet in 1990 and 1991. Between 1988 and 1991, he served as senior military assistant to Secretaries of Defense Frank Carlucci and Dick Cheney, the senior military position in the Office of the Secretary of Defense. In 1988, the admiral was the director of the Office of Program Appraisal for the secretary of the Navy. In 1987, he served as commander of Submarine Group Six, the Navy's largest submarine group, with 20 strategic ballistic missile

submarines, 45 nuclear attack submarines, and more than 15,000 men and women. Earlier in his career, he commanded Submarine Squadron Four, the USS Sam Houston, and the USS City of Corpus Christi. Admiral Owens has written more than 50 articles on national security and wrote the book *High Seas*. His latest book, *Lifting the Fog of War*, was published in April 2000 and revised and republished in 2008. He is a 1962 graduate of the U.S. Naval Academy and holds a B.S. in mathematics. He also holds bachelor's and master's degrees in politics, philosophy, and economics from Oxford University and a master's degree in management from George Washington University. The admiral is the founder of Extend America, a 5-year state wireless telecommunications venture, and also sits on the public boards of Polycom, Wipro, and Daimler AG as well as the private boards of Intelius, Force 10 Networks, Unifrax, and AEA Investors LLC. Owens is a member of several philanthropic boards including the Carnegie Foundation, the Brookings Institution, and the Fred Hutchinson Cancer Research Center. He is also a member of the Canadian Council of Chief Executives and the Council on Foreign Affairs.

Kenneth W. Dam, *Co-chair*, University of Chicago, has devoted his career to public policy issues, both as a practitioner and as a professor. He served as deputy secretary (the second-ranking official) in the Department of Treasury (2001-2003) and in the Department of State (1982-1985). In 1973 he was executive director of the Council on Economic Policy, a White House office responsible for coordinating U.S. domestic and international economic policy. From 1971 to 1973 he served as assistant director for national security and international policy at the Office of Management and Budget. He began his Washington career as law clerk to U.S. Supreme Court Justice Charles E. Whittaker (1957-1958). Professor Dam's entire academic career has been devoted to the University of Chicago, beginning in 1960 and extending, with various leaves of absence, to the present. From 1980 to 1982 he served as provost of the University of Chicago. Most of his academic work has centered on law and economics, particularly with respect to international issues. Professor Dam's other activities include serving as IBM vice president for law and external relations (1985-1992) and as president and chief executive officer of the United Way of America for a 6-month period in 1992. He has extensive experience as an arbitrator. The professor is a member of the board of the Brookings Institution and serves as a senior fellow of that organization. He is a member of the Shadow Financial Regulatory Committee and of the National Research Council's Science, Technology and Law Panel. He has been elected to membership in the American Law Institute and the American Academy of Arts and Sciences. He was chair of the German-American Academic Council and a board member of a number of non-profit institutions, including the Council on Foreign Relations (New

York) and the Chicago Council on Foreign Relations. He currently serves on the board of the Financial Services Volunteer Corps. Professor Dam served for 13 years on the board of Alcoa. He received a B.S. in 1954 from the University of Kansas, a J.D. in 1957 from the University of Chicago, and an LL.D. (hon.) in 1983 from the New School for Social Research. The professor served as chair for the CSTB committee that produced the report *Cryptography's Role in Securing the Information Society*, and he served on the CSTB committee that produced the report *Global Networks and Local Values: A Comparative Look at Germany and the United States*.

Thomas A. Berson, president of Anagram Laboratories, has spent his career working both the defensive and the offensive sides of the information security battle. After stints as a researcher, a cold warrior, and Silicon Valley entrepreneur, Dr. Berson founded Anagram Laboratories, a thriving information security consultancy that is celebrating its 23rd anniversary in 2009. He is attracted most strongly to security issues raised at the confluence of technology, business, and world events. His client base includes Salesforce.com (disruptive at the center of the net) and Skype (disruptive at the edge). Dr. Berson is a student of Sun Tzu's Art of War and its applicability to modern information conflict. Dr. Berson was the first person to be named a fellow of the International Association for Cryptologic Research. His citation reads, "for visionary and essential service and for numerous valuable contributions to the technical, social, and commercial development of cryptology and security." Dr. Berson was an editor of the *Journal of Cryptology* for 14 years. He is a past chair of the IEEE Technical Committee on Security and Privacy and a past president of the International Association for Cryptologic Research. He earned a B.S. in physics from the State University of New York in 1967 and a Ph.D. in computer science from the University of London in 1977. He was a visiting fellow in mathematics at the University of Cambridge and is a life member of Clare Hall, Cambridge. Dr. Berson has been a member of two previous National Research Council committees: the Committee on Computer Security in the Department of Energy and the Committee to Review DOD C4I Plans and Programs.

Gerhard Casper is president emeritus of Stanford University and the Peter and Helen Bing Professor in Undergraduate Education at Stanford. He is also a professor of law, a senior fellow at the Freeman Spogli Institute for International Studies, and a professor of political science (by courtesy). Mr. Casper studied law at the Universities of Freiburg and Hamburg, where, in 1961, he earned his first law degree. He attended Yale Law School, obtaining a master of laws degree in 1962. He then returned to Freiburg, where he received his doctorate in 1964. He has been awarded honorary doctorates, most recently in law from Yale and in philosophy from Uppsala. In the fall of 1964, Mr. Casper emigrated to

the United States, spending 2 years as an assistant professor of political science at the University of California at Berkeley. In 1966, he joined the faculty of the University of Chicago Law School, and between 1979 and 1987 served as its dean. In 1989, Mr. Casper was appointed provost of the University of Chicago. He served as president of Stanford University from 1992 to 2000. Mr. Casper has written and taught primarily in the fields of constitutional law, constitutional history, comparative law, and jurisprudence. From 1977 to 1991, he was an editor of the *Supreme Court Review*. His books include a monograph on legal realism (Berlin, 1967), an empirical study of the Supreme Court's workload (Chicago, 1976, with Richard A. Posner), and *Separating Power* (Cambridge, Mass., 1997), concerning the separation of powers practices at the end of the 18th century in the United States. About the Stanford presidency, he wrote *Cares of the University* (Stanford, 1997). He is also the author of numerous scholarly articles and occasional pieces. He has been elected to the American Law Institute (1977), the International Academy of Comparative Law, the American Academy of Arts and Sciences (1980), the Ordre pour le mérite for Sciences and the Arts (1993), and the American Philosophical Society (1996). At present, Mr. Casper serves as a member of the board of trustees of the Central European University in Budapest as well as a member of the board of trustees of the American Academy in Berlin. He is also a member other boards, including the Council of the American Law Institute and the Committee for Economic Development. From 1998 to 2005, he was a member of the Trilateral Commission and, from 2000 to 2008, he served as a successor trustee of Yale University.

David D. Clark, NAE, has worked at the Massachusetts Institute of Technology's (MIT's) Computer Science and Artificial Intelligence Laboratory, where he is currently a senior research scientist in charge of the Advanced Network Architecture Group, since receiving his Ph.D. from MIT in 1973. Dr. Clark's research interests include networks, network protocols, operating systems, distributed systems, and computer and communications security. After receiving his Ph.D., he worked on the early stages of the ARPANET and on the development of token ring local area network technology. Since the mid-1970s, Dr. Clark has been involved in the development of the Internet. From 1981 to 1989, he acted as chief protocol architect in this development and chaired the Internet Activities Board. His current research looks at redefinition of the architectural underpinnings of the Internet, and the relation of technology and architecture to economic, societal, and policy considerations. He is helping the U.S. National Science Foundation organize its Future Internet Design program. In the security area, Dr. Clark participated in the early development of the multilevel secure Multics operating system. He developed an information security model that stresses integrity of data rather than

disclosure control. Dr. Clark is a fellow of the Association for Computing Machinery and the IEEE and is a member of the National Academy of Engineering. He received the ACM SIGCOMM award and the IEEE award in international communications, as well as the IEEE Hamming Award for his work on the Internet. He is a consultant to a number of companies and has served on a number of technical advisory boards. Dr. Clark was the past chair of the Computer Science and Telecommunications Board (CSTB) at the National Research Council. He chaired the committee that produced the CSTB report *Computers at Risk: Safe Computing in the Information Age*. Dr. Clark also served on the committees that produced the CSTB reports *Toward a National Research Network, Realizing the Information Future: The Internet and Beyond*, and *The Unpredictable Certainty: Information Infrastructure Through 2000*. Dr. Clark graduated from Swarthmore College in 1966 and received a Ph.D. from MIT in 1973.

Richard L. Garwin, NAS/NAE/IOM, is an IBM fellow emeritus at the Thomas J. Watson Research Center and an adjunct professor of physics at Columbia University. Dr. Garwin is a physicist with expertise in intelligence and in nuclear, chemical, and biological weapons and defenses. From 1994 to 2001 he chaired the Arms Control and Nonproliferation Advisory Board at the Department of State. Dr. Garwin received the Enrico Fermi Award of the President and the Department of Energy (1996) and the R.V. Jones Intelligence Award of the U.S. government intelligence community (1996). In 2003 he received the National Medal of Science and in 2000 was named by the National Reconnaissance Office as one of its 10 founders of national reconnaissance. Dr. Garwin's publications include *Megawatts and Megatons: The Future of Nuclear Power and Nuclear Weapons* (2003); *Megawatts and Megatons: A Turning Point for the Nuclear Age?* (2001); *Control of Nuclear Arms at Crossroads* (2000); *A Defense That Will Not Defend* (2000); *Boost-Phase Intercept: A Better Alternative* (2000); *Feux Follets et Champignons Nucléaires* (1997); and *Management and Disposition of Excess Weapons Plutonium* (1994). Dr. Garwin has a Ph.D. and an M.S. in physics from the University of Chicago (1949, 1948) and a B.S. in physics from Case Western Reserve University (1947). He has never been a member of any private boards. Many of his papers and much testimony is posted at <http://www.fas.org/RLG/>.

Jack L. Goldsmith III has been a professor of law at Harvard Law School since 2004. In 2003-2004 he was the assistant attorney general in the U.S. Department of Justice's Office of Legal Counsel. At that time he was also a professor of law at the University of Virginia Law School. Before that he served on the faculty of the University of Chicago Law School and as special counsel to the General Counsel in the Department of Defense. Earlier Mr. Goldsmith was an associate professor at the University of Virginia Law School from 1994 to 1997. Mr. Goldsmith received a B.A.

in philosophy *summa cum laude* from Washington and Lee University in 1984, a B.A. in philosophy, politics, and economics from Oxford University in 1986, a J.D. from Yale Law School in 1989, and a diploma in private international law from The Hague Academy of International Law in 1992. After law school he clerked for Judge J. Harvie Wilkinson of the United States Court of Appeals for the Fourth Circuit, Justice Anthony M. Kennedy of the Supreme Court of the United States, and Judge George A. Aldrich of the Iran-U.S. Claims Tribunal. He also previously has served as an associate at Covington & Burling. Mr. Goldsmith's scholarly interests include international law, foreign relations law, national security law, conflict of laws, and civil procedure.

Carl G. O'Berry is with the Boeing Company, where he is vice president of Network-Centric Architectures. He retired from the U.S. Air Force as a lieutenant general in August 1995. Until December 1998 he was vice president and director of planning and information technology for the Space and Systems Technology Group at Motorola, where he was responsible for groupwide strategic and long-range planning and executive management of group information technology solutions and services. In addition, he was responsible for information technology architectures and road maps, new information technology business development, and leadership of information technology innovation and process reengineering. He was previously deputy chief of staff for Command, Control, Communications and Computers at U.S. Air Force headquarters, a position from which he directed Air Force-wide information systems planning and policy development. Earlier in his Air Force career, he served as commander of the Air Force Rome Air Development Center and as joint program manager of the World-Wide Military Command and Control System Information System. He also led the development and field testing of an airborne radar sensing/tracking system that was the forerunner of the Joint Surveillance and Target Attack Radar System. He has a master's degree in systems management from the Air Force Institute of Technology and a bachelor's degree in electrical engineering from New Mexico State University. He served on the NRC committee that produced *Realizing the Potential of C4I: Fundamental Challenges*.

Jerome H. Saltzer, NAE, is a professor of computer science, emeritus, in the Department of Electrical Engineering and Computer Science at MIT. A member of that department since 1961, he helped formulate the original undergraduate curriculum in computer science and led the development of the core subject on the engineering of computer systems. At the MIT Computer Science and Artificial Intelligence Laboratory he designed one of the earliest widely used word-processing systems; he participated in the development of the Multics system, for which he designed the kernel thread package and with students and colleagues developed the security

mechanisms and what would today be known as a microkernel; and together with David Clark and David Reed, he articulated the end-to-end argument, a key organizing principle of the Internet. He was also involved in the design of a token-passing ring local area network, the networking of personal computers, the Kerberos single-login authentication system, and digital library systems. Dr. Saltzer was technical director of MIT Project Athena, a system for undergraduate education and an early example of a system organization now called "cloud computing." Throughout his work, he has had a particular interest in the impact of computer systems on privacy and the risks of depending on fragile technology. Dr. Saltzer is a fellow of the IEEE and the AAAS; a member of the Association for Computing Machinery, the ACM Committee on Computers and Public Policy, and the Catalog Raisonné Scholars Association; a former member of the Computer Science and Telecommunications Board of the National Research Council; and a former member of the mayor's Telecommunications Advisory Board for the City of Newton, Massachusetts. Dr. Saltzer received an S.B. (1961), an S.M. (1963), and an Sc.D. (1966), from MIT, all in the field of electrical engineering.

Mark Seiden is a consultant with MSB Associates. Previously he was a senior consultant with Cutter's Business-IT Strategies Practice and a member of the Leadership Group of the Cutter Consortium's Risk Management Intelligence Network. He has consulted since 1983 in the areas of security, network, and software engineering to companies worldwide, with clients including start-ups, major computer and communication companies, financial institutions, law firms, UN agencies, online content providers, ISPs, research organizations, and non-profits. As an independent consultant and in varying roles at Securify (also known as the Kroll O'Gara Information Security Group), his most recent projects have included design, architecture, and implementation for e-business systems; security for online financial transaction processing and distributed document-processing systems; custom firewalls based on open-source components; finding computer criminals; and penetration testing the network and physical security of deployed systems, enterprises, and collocation facilities. Mr. Seiden has 35 years' programming experience. He has been a Unix and mainframe system programmer; written Macintosh applications; spent time at IBM Research, Xerox Parc, Bell Labs, and Bellcore; and has taught at the university level. Mr. Seiden has been on the board of directors of two user groups and is on the Technical Advisory Board of Counterpane Security Systems. Mr. Seiden has an M.S. in computer science/electrical engineering from Columbia University and as an undergraduate at Columbia studied math, music, and linguistics.

Sarah Sewall is the director of the Carr Center at the John F. Kennedy School of Government at Harvard University and lecturer in public policy,

and she also directs the Carr Center's Program on National Security and Human Rights. During the Clinton administration, Ms. Sewall served in the Department of Defense as the first deputy assistant secretary for Peacekeeping and Humanitarian Assistance. From 1987 to 1993, she served as senior foreign policy adviser to Senate Majority Leader George J. Mitchell, was a delegate to the Senate Arms Control Observer Group, and was on the Senate Democratic Policy Committee. Ms. Sewall has also worked at a variety of defense research organizations and as associate director of the Committee on International Security Studies at the American Academy of Arts and Sciences. She was lead editor of *The United States and the International Criminal Court: National Security and International Law* (2000) and has written widely on U.S. foreign policy, multilateralism, peace operations, and military intervention. Her current research focuses on the civilian in war and includes facilitating a dialogue between the military and human rights communities on the use of force.

Walter B. Slocombe practices in Caplin & Drysdale's office in Washington, D.C. He served as undersecretary of defense for policy from 1994 to 2001, and as senior advisor for national defense in the Coalition Provisional Authority for Iraq in 2003. In 2004, President Bush appointed him to the Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction. He served on the National Security Council staff in 1969-1970, and as principal deputy assistant secretary of defense for international security affairs in 1977-1979 and deputy undersecretary for policy in 1979-1981. He has also been a member of various advisory or governing boards of several academic and defense analysis institutions and government agencies, including a panel of the National Security Agency's advisory board. Mr. Slocombe was awarded the Department of Defense's Distinguished Public Service Medal (1981, 1995, 1997, 2001, 2004) and its Joseph Kruzal Award for Distinguished Service in the Pursuit of Peace (2000) and has been named an honorary submariner by the Fleet Submarine Force. His international service has been recognized by awards from the Polish, German, and Korean governments. Mr. Slocombe has published numerous articles and monographs on tax law issues and on defense policy and organization. He is a 1963 graduate of Princeton University, attended Balliol College Oxford as a Rhodes Scholar, and in 1968 received his law degree from Harvard Law School where he was note editor of the *Law Review*.

William O. Studeman retired from the U.S. Navy in 1995 with the rank of admiral. A top-level military manager and government leader, his flag positions included director of the Navy Long-Range Planning Group and executive secretary of the Advanced Technology Panel of the CNO Executive Board, director of naval intelligence, and director of the National Security Agency. In 1992, President Bush nominated him to

the political position of deputy director of Central Intelligence. Between 1992 and 1995, Mr. Studeman served as deputy to Robert Gates, James Woolsey, and John Deutch and served twice for extended periods as the acting director of Central Intelligence. In this capacity, he was the intelligence community's representative to the President's Management Council and responsible for implementing the National Performance Review for downsizing, streamlining, and reengineering the federal government. He has conducted extensive operational intelligence tours overseas. Some of his key tours included duty as executive assistant to both the director of naval intelligence and the vice chief of naval operations; officer in charge of the Atlantic Fleet Ocean Surveillance Information Center; commanding officer of the Navy Operational Intelligence Center, and assistant chief of staff for intelligence, U.S. Sixth Fleet staff at Gaeta, Italy. In addition to his management and ISR experience, he has extensive background in anti-submarine warfare, C4ISR, information warfare, and homeland security. In 2005, he retired from Northrop Grumman Mission Systems, where he was sector vice president and deputy general manager for intelligence and information superiority, and where he also coordinated the sector's homeland security activities and technology partnerships. Before joining TRW (which was acquired by Northrop Grumman in December 2002) in September 1996, Mr. Studeman worked for a year consulting on defense, intelligence, information infrastructure, security, and management issues, following 34 years of career military service. He is a distinguished graduate of the Defense Intelligence School, the Naval War College, and the National War College. He received a bachelor's degree in history from the University of the South, a master's degree in public and international affairs from George Washington University, and numerous honorary degrees. Mr. Studeman also serves on numerous government boards, including the Defense Science Board and the Presidential Commission on WMD.

Michael A. Vatis is a partner in the New York office of Steptoe & Johnson LLP. His practice focuses on the Internet, e-commerce, and technology matters, with special emphasis on issues involving security, intelligence, and law enforcement. He also is an experienced appellate litigator. Mr. Vatis has spent most of his career addressing cutting-edge issues at the intersection of law, policy, and technology. He was the founding director of the National Infrastructure Protection Center at the FBI, the first government organization responsible for detecting, warning of, and responding to cyberattacks, including computer crimes, cyberterrorism, cyber-espionage, and information warfare. Before that, Mr. Vatis served as associate deputy attorney general and deputy director of the Executive Office for National Security in the Department of Justice, where he advised the attorney general and deputy attorney general and coordi-

nated the department's activities involving counterterrorism, intelligence, and cybercrime. In that capacity, he also helped lead the development of the nation's first policies on critical infrastructure protection. Mr. Vatis served as special counsel at the Department of Defense, where he handled sensitive legal and policy issues for the secretary and deputy secretary of defense and the general counsel, receiving the Secretary of Defense Award for Excellence. After leaving the government in 2001, Mr. Vatis served as the first director of the Institute for Security Technology Studies at Dartmouth, a federally funded counterterrorism and cybersecurity research institute. He was simultaneously the founding chairman of the Institute for Information Infrastructure Protection (I3P). I3P, a consortium of leading cybersecurity research organizations, worked with industry, government, and academia to develop a comprehensive research and development agenda to improve the security of the nation's computer and communications networks. Mr. Vatis also served as the executive director of the Markle Task Force on National Security in the Information Age, a highly influential group of technology company executives, former government officials, and civil libertarians that examined how the government could more effectively use information and technology to combat terrorism while preserving civil liberties. He was the principal author of the group's second report, whose recommendations were adopted by the 9/11 Commission and included in the 2004 Intelligence Reform Act. Mr. Vatis has regularly testified before congressional committees on counterterrorism, intelligence, and cybersecurity issues. He is also interviewed on television, radio, and in print media and has been a guest lecturer at many prestigious law schools and universities and a frequent speaker at industry conferences worldwide.

STAFF MEMBERS

Herbert S. Lin, the study director, is chief scientist for the National Research Council's Computer Science and Telecommunications Board, where he has been a study director for major projects on public policy and information technology. These studies include a 1996 study on national cryptography policy (*Cryptography's Role in Securing the Information Society*), a 1991 study on the future of computer science (*Computing the Future*), a 1999 study of Defense Department systems for command, control, communications, computing, and intelligence (*Realizing the Potential of C4I: Fundamental Challenges*), a 2000 study on workforce issues in high technology (*Building a Workforce for the Information Economy*), a 2002 study on protecting kids from Internet pornography and sexual exploitation (*Youth, Pornography, and the Internet*), a 2004 study on aspects of the FBI's information technology modernization program (*A Review of the FBI's Tril-*

ogy IT Modernization Program), a 2005 study on electronic voting (*Asking the Right Questions About Electronic Voting*), a 2005 study on computational biology (*Catalyzing Inquiry at the Interface of Computing and Biology*), a 2007 study on privacy and information technology (*Engaging Privacy and Information Technology in a Digital Age*), a 2007 study on cybersecurity research (*Toward a Safer and More Secure Cyberspace*), and a 2009 study on health care information technology (*Computational Technology for Effective Health Care*). Before his NRC service, he was a professional staff member and staff scientist for the House Armed Services Committee (1986-1990), where his portfolio included defense policy and arms control issues. He received his doctorate in physics from MIT. Apart from his CSTB work, he is published in cognitive science, science education, biophysics, and arms control and defense policy. He also consults on K-12 math and science education.

Ted Schmitt was a consultant for the Computer Science and Telecommunications Board of the National Research Council until 2008. He was involved in the CSTB projects on offensive information warfare, biometrics, and wireless technology. Recently completed projects he worked on include a review of health IT standards efforts at the Office of the National Coordinator for Health IT, a comprehensive exploration of cybersecurity and the use of IT to enhance disaster management. Before joining CSTB, Mr. Schmitt was involved in the development of the digital media industry and played an active role in various related media standards groups. Prior to that, he served as technical director at a number of small technology companies in Germany, Sweden, and the United States. He started his career in 1984 as a software engineer for IBM, earning two patents and several technical achievement awards. Mr. Schmitt received an M.A. in international science and technology policy from George Washington University. He received a B.S. in electrical engineering in 1984 and a B.A. in German in 1997 from Purdue University, and he studied at the University of Hamburg, Germany.

Appendix B

Meeting Participants and Other Contributors

The Committee on Offensive Information Warfare held five open meetings starting in June 2006. These meetings included information-gathering sessions open to the public, as well as closed segments for committee deliberation. The committee heard from numerous presenters at these meetings, including the following.

MEETING 1, JUNE 26-27, 2006

Thomas Wingfield, Potomac Institute
Rod Wallace, Nortel
Steven Bellovin, Columbia University

MEETING 2, OCTOBER 30-31, 2006

K.A. Taipale, Center for Advanced Studies in Science and Technology
Policy
Stuart Starr, Center for Technology and National Security Policy,
National Defense University
William Howard, Independent Consultant
Linton Wells, Department of Defense
Thomas Schelling, University of Maryland (videotaped)
LTC Eric Jensen, Office of the Judge Advocate General, U.S. Army
Lt. Gen. Bill Donahue, U.S. Air Force (retired)
Joe Dhillon, University of the Pacific, McGeorge School of Law

Neal Pollard, Georgetown University and NCTC
Admiral Elizabeth Hight, Joint Task Force on Global Network
Operations
Jeff McMahan, Rutgers University
Father J. Bryan Hehir, Harvard University/Catholic Charities

FACT-GATHERING SESSION, JANUARY 27, 2007

Forrest Hare, U.S. Air Force

MEETING 3, FEBRUARY 20-21, 2007

Sam Gardiner, U.S. Air Force (retired)
James N. Miller, Jr., Hicks and Associates, Inc.
Dorothy E. Denning, Naval Postgraduate School
Dan Kuehl, National Defense University
Jeff Smith, former CIA General Counsel

FACT-GATHERING SESSION, MARCH 28, 2007

Stephen A. Cambone, former Undersecretary of Defense for Intelligence

MEETING 4, APRIL 10-11, 2007

Patrick D. Allen, General Dynamics Advanced Information Systems
David Koplou, Georgetown University
Milo Medin, M2Z Networks
Jeffrey I. Schiller, MIT
Jim Dempsey, Center for Democracy and Technology
Richard Salgado, Yahoo!
Eugene Volokh, UCLA School of Law
Robert Weisberg, Stanford Law School
Helen Stacy, Stanford University
Naveen Jain, Intelius

Appendix C

Illustrative Criminal Cyberattacks

THE INVITA CASE

In 2001, the FBI arrested two Russians, Alexey Ivanov, 21, and Vasily Gorshkov, 25, who were accused of breaking into dozens of sites ranging from Internet service providers to banks.¹ Where they found financial records they could steal, they stole financial records. Where they couldn't, they contacted the sites saying they knew about a recent break-in and offered their services to remediate the problems or they threatened to release other information stolen from the site to damage the victim's public reputation. The FBI took advantage of the solicitations for work to lure the two suspects to the United States on the pretext of a job interview, where the interviewees were arrested. Approximately 2.3 gigabytes (compressed) of evidentiary data was remotely seized from the suspects' server in Russia before it was taken offline by others still in Russia. Both were convicted in separate U.S. district courts. Gorshkov was charged with damages in excess of \$2.5 million and ordered to both serve jail time and pay a combined total of nearly \$1.5 million in restitution.

When analyzed, the evidence—lists of credit cards numbers, Perl scripts for manipulating e-mail and auction accounts, and other hacking tools—showed a complex scheme involving the creation of fake anonymous e-mail accounts and fake eBay seller and PayPal customer accounts, all fueled by the stolen financial information they possessed. They would

¹Department of Justice, "Russian Computer Hacker Sentenced to Three Years in Prison," 2002, available at <http://www.usdoj.gov/criminal/cybercrime/gorshkovSent.htm>.

create a fake auction item with a value less than \$500 to avoid triggering fraud alarms. They would use other fake accounts to bid on the item, and they knew how to rig the bidding so they would always win (thus not defrauding any real bidders who might report the activity). The fake PayPal accounts would be used to clear the transaction, and they even used the fake bidder accounts to “rate the seller,” inflating the credibility of the fake accounts.

One very interesting aspect of this case is the automation of all processes related to e-mail account creation and management, online payment account creation and management, web-based transaction processing, and electronic funds transfer. Tens of thousands of stolen credit card numbers were carefully used in ways that limited the losses to less than a few hundred dollars per card. The automation allowed the group to focus on the intrusions, data exfiltration and sorting, and other aspects of their activity that brought in money. This was all done by a small group of perhaps a half-dozen individuals,² skilled programmers who could not find jobs locally that paid anything near what their skills were worth. Ivanov was described by U.S. District Court Judge Thompson as a “manager or supervisor,” while Gorshkov claimed he was “the boss.” (Both statements could be true if there are six or more individuals involved.) They claim to have worked up to 16 hours per day over about 1 year³ and to have generated \$150,000 in 6 months. This is enough to pay the salaries of 20 (unemployed) Russian rocket scientists at 2003 salary rates.⁴

THE ISRAELI TROJAN HORSE INDUSTRIAL ESPIONAGE CASE

In 2005, a couple were arrested in Britain on charges of creating a Trojan horse key logger and installing it on systems at dozens of sites by way of CD-ROMs containing what was purported to be a business proposal.⁵ This has been described as the largest industrial espionage case in Israeli history. The espionage activity was primarily targeted at competitors to the clients of three private investigation firms, at a cost

² Philip Attfield, “United States v Gorshkov Detailed Forensics and Case Study; Expert Witness Perspective,” in *Proceedings of the First International Workshop on Systematic Approaches to Digital Forensic Engineering (SADFE05)*, 2005, available at <http://ieeexplore.ieee.org/iel5/10612/33521/01592518.pdf?arnumber=1592518>.

³ Art Jahnke, “Russian Roulette,” 2005, available at <http://www.csoonline.com/read/010105/russian.html>.

⁴ Stephanie Overby, “Big Ideas 2003: Passages Beyond India,” 2003, available at http://www.cio.com/article/31589/Big_Ideas_Passages_Beyond_India/1.

⁵ See, for example, Avi Cohen, “Scandal Shocks Business World,” 2005, available at <http://www.ynetnews.com/articles/0,7340,L-3091900,00.html>. See also Bob Sullivan, “Israel Espionage Case Points to New Net Threat,” June 9, 2005, available at <http://www.msnbc.msn.com/id/8145520/>.

of approximately \$4,000 per compromised computer. Eighteen people were arrested and questioned in the case; however, it was primarily just a couple and their 17-year-old son who were responsible for software production, distribution, and data collection services. It was reported that about a hundred pieces of computer equipment were seized by authorities at the time of arrest. The espionage activity was believed to have gone on for a year and a half, partly because the Trojan was highly targeted. The suspects were identified because of a personal vendetta having to do with a bitter divorce trial, and not because they were detected in the acts of computer intrusion or data exfiltration from the corporate victims.

In this case, the goal was to compromise the confidentiality of business records by means of unauthorized access and data exfiltration from compromised computers. The 100 items of equipment seized by authorities were probably development hosts, file servers that received exfiltrated files, and perhaps processing hosts that would assist in sifting through the files collected by the Trojan horse malware. It is not publicly known how sophisticated the operation was, but the number of arrests suggests that a significant amount of high-level intellectual property theft had taken place as part of this operation.

OPERATIONS “CYBERSLAM,” “BOTMASTER UNDERGROUND,” AND OTHER BOTNET CASES

The computer security news media are full of stories of botnets—huge numbers of compromised personal computers running Internet Relay Chat (IRC) robot programs, or “bots” for short⁶—being used to automate many types of criminal activity, from delivery of spam, to theft of software license keys, to distributed denial-of-service (DDOS) attacks for extortion or other financial gain, to click fraud. Four prominent incidents that received attention were these:

- In one of the first cases of DDOS-for-hire, Saad “Jay” Echouafni, the owner of a satellite TV equipment sales company, hired someone known for running large DDOS attack botnets, paying him or her \$150,000 per year. This person, in turn, subcontracted the work to four other individuals who managed their own botnets. The purpose was to carry out extended DDOS attacks against Echouafni’s business competitors. Specific new attack mechanisms were coded into Agobot, the bot software being used by several of the subcontractors, in order to defeat DDOS

⁶ For a description of bots and botnets, see “What Is a Botnet?,” available at <http://www.techfaq.com/botnet.shtml>.

mitigation techniques employed to protect the targeted sites. The result was an estimated \$2 million in lost revenue and cost of cleanup.⁷

- Jeanson James Ancheta entered a plea of guilty to taking control of approximately 400,000 personal computers (including computers at the Naval Warfare Center at China Lake and the Defense Information Systems Agency in Virginia) for criminal purposes, including selling access to DDOS botnets and performing click fraud. Ancheta maintained a series of servers that coordinated the bot activity, including operating private channels for command and control of the bots that were sold to third parties wishing to use them for their own criminal purposes (e.g., denial of service attacks and spam transmission), as well as for supporting these “customers.” He admitted to collecting more than \$107,000 in advertising affiliate proceeds from directing the bots on compromised computers into referring him and another unindicted co-conspirator to the adware sites (known as “click fraud.”) The income from these operations funded the servers and hosting costs and allowed Ancheta to purchase a new BMW with cash, all of which was returned as part of the plea agreement.⁸

- Prosecutors in the Netherlands stated publicly that they believe three teenage suspects, two of whom were convicted and sentenced in February 2007, controlled as many as 1.5 million personal computers worldwide using a variant of the ToxBot program. The three were accused of using these botnets to steal credit card numbers and other personal data and to blackmail online businesses.⁹

- In June 2007, the FBI reported an event of similar size in the United States, part of “Operation Bot Roast,” involving over 1 million personal computers. Arrested were three individuals, two accused of performing DDOS attacks and one reported to be one of the most prolific spammers at the time.¹⁰

In all of these cases, small groups of relatively young people with skills in programming and computer system administration were able to successfully compromise and control over a million personal comput-

⁷ Department of Justice, “Criminal Complaint: United States of America v. Paul G. Ashley, Jonathan David Hall, Joshua James Schichtel, Richard Roby and Lee Graham Walker,” 2004, available at <http://www.reverse.net/operationcyberslam.pdf>.

⁸ Department of Justice, “Computer Virus Broker Arrested for Selling Armies of Infected Computers to Hackers and Spammers,” 2005, available at <http://www.cybercrime.gov/anchetaArrest.htm>.

⁹ Joris Evers, “Bot Herders’ May Have Controlled 1.5 million PCs,” 2005, available at <http://news.com.com/Bot+herders+may+have+controlled+1.5+million+PCs/2100-73503-5906896.html>.

¹⁰ Department of Justice, “Over One Million Potential Victims of Botnet Cyber Crime,” 2007, available at <http://www.ic3.gov/media/initiatives/BotRoast.pdf>.

ers around the world, using very little additional software above and beyond modified versions of publicly available IRC-based botnet and IRC server software. These are just the proverbial tip of the iceberg in terms of online crime using distributed intruder tool networks, including botnets. A migration is beginning to take place, away from the easier to detect and mitigate IRC botnets and toward the use of heavily encrypted peer-to-peer malicious programs for distributed command and control.

THE STAKKATO INTRUSIONS

In 2003, a teenager in Sweden began a series of intrusions that lasted through 2005 and compromised more than 1000 hosts at supercomputer centers, national labs, universities, corporations, and military bases around the world.¹¹ The initial target of attack was remotely exploitable vulnerabilities in Linux systems, where a rootkit named SuckKIT was installed that hides itself on the system and logs all keystrokes. This allowed the attacker to steal account/password credentials of people logging into the compromised host or using that host to log in to some other host (possibly at another site). The attacker would sometimes replace the login message with a taunt about how using Linux was a great way to share accounts.

One aspect of the Stakkato case that is not appreciated by many is the clever exploitation of the implicit trust relationships that exist between systems based on users having accounts on more than one system, and more than one user sharing any given system. The attacker would steal passwords to gain access to accounts, and then do sufficient mapping of login relationships between hosts to infer where these same login/password combinations might work. He would then log into those systems, preferably using administrator accounts, and then repeat the process of installing the keystroke logger and further extending his reach into new systems and networks: (1) University researchers often have appointments in multiple institutions, or multiple departments within an institution; (2) those researchers have contractual relationships with corporations in industry; (3) supercomputer centers are used by researchers in academia, in business, and in the military; (4) the same business that employs a researcher in one field (who may require the services of a supercomputer center) may also be involved in software or hardware engineering and sales. Stakkato probably did not even plan on it, but during the compromise of those 1000+ systems, an account at Cisco Systems was compromised and was used to obtain a copy of part of the Cisco IOS router software base, which was later posted on a Russian website. The

¹¹ Leif Nixon, "The Stakkato Intrusions," 2006, available at <http://www.nsc.liu.se/nixon/stakkato.pdf>.

nature of the login trust relationships between sites was one reason the intrusions lasted so long: Some sites would clean up their systems, only to find them compromised again a short time later because they did not realize the extent of shared access between systems, nor did they realize what the compromise of passwords through keystroke logging means in terms of completely mitigating an attack of this nature.

TJX FINANCIAL DATA THEFTS

At various dates between July 2005 and January 2006, intruders used access to systems within the corporate network of TJX Companies, Inc., to obtain and exfiltrate 45.7 million payment card (i.e., credit or debit card) records.¹²

In March 2007, six suspects were arrested, with four more at large, all believed to be involved in the data theft and an elaborate scheme for using the stolen data to make an estimated \$8 million in purchases of gift cards and electronics equipment.¹³ This is on par with the number of individuals involved in the Invita case, the first case in this appendix. However the financial damage involved in the TJX case could be orders of magnitude greater than the losses in the Invita case just 5 years earlier. Based on estimates of \$50 to \$250 per record, the TJX breach could cost the company in excess of \$2 billion. Several pending lawsuits and a regulatory investigation are also underway.

As of the time of this writing, few details about the attack mechanism have been made public, but it would be reasonable to assume an attack methodology similar to that in the previous cases. Since the attackers were in the networks for over a year, there was a great deal of time available to quietly exploit stolen credentials and explore the network, identifying the crown jewels in terms of financial information databases.

¹² The SEC Form 10-K filing by TJX claims that, in general, track 2 data—all data, including the PIN number on debit cards, necessary to clone the card—was either masked off with asterisks or stored in encrypted form. TJX does, however, state that, “despite our masking and encryption practices on our Framingham system in 2006, the technology utilized in the Computer Intrusion during 2006 could have enabled the Intruder to steal payment card data from our Framingham system during the payment card issuers’ approval process, in which data (including the track 2 data) is transmitted to payment card issuers without encryption. Further, we believe that the Intruder had access to the decryption tool for the encryption software utilized by TJX.” This means there is a possibility that payment cards could be cloned by the attackers.

¹³ Jenn Abelson, “Breach of Data at TJX Is Called the Biggest Ever: Stolen Numbers Put at 45.7 Million,” March 29, 2007, available at http://www.boston.com/business/globe/articles/2007/03/29/breach_of_data_at_tjx_is_called_the_biggest_ever/.

Appendix D

Views on the Use of Force in Cyberspace

COMPUTER NETWORK ATTACK AND THE USE OF FORCE IN INTERNATIONAL LAW

In 1999, Michael Schmitt addressed the issue of cyberattack as a use of force.¹ Focusing on computer network attack (CNA) (remote-access attack, as described in Chapter 2), Schmitt argued that CNA should be understood in terms of its effects and said that the consequences of a CNA rather than its specific modality were the most important factor in its categorization. He focused on the consequences of a CNA because of their potentially broad range: “CNA spans the spectrum of consequentiality. Its effects freely range from mere inconvenience (e.g., shutting down an academic network temporarily) to physical destruction (e.g., as in creating a hammering phenomenon in oil pipelines so as to cause them to burst) to death (e.g., shutting down power to a hospital with no backup generators).”

Thus, for example, Schmitt argued that “CNA specifically intended to directly cause physical damage to tangible property or injury or death to human beings is reasonably characterized as a use of armed force,” and so “pipeline destruction and the shutting of power to the hospital are examples of CNA which the actor knows can, and intends to, directly cause destruction and serious injury.” He further noted that “armed coercion

¹ Michael Schmitt, “Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework,” *Columbia Journal of Transnational Law* 37:885-937, 1999.

is not defined by whether or not kinetic energy is employed or released, but rather by the nature of the direct results caused, specifically physical damage and human injury.”

On the other hand, Schmitt noted that economic coercion is not generally regarded as rising to the level of a “use of force,” so that a CNA that seeks economic coercion cannot be considered a use of force. For a CNA to be considered a use of force, he argued that it must be more consequential than simple economic coercion but does not necessarily have to meet the threshold of being considered a use of “armed force” as described in the previous paragraph. He thus argues that “the use of force line must lie somewhere between economic coercion and the use of armed force.”

Schmitt then offered a seven-element framework for categorizing computer network attack as a use of force:

- *Severity.* If people are killed or there is extensive property damage, the action is probably military; the less damage, the less likely the action is a use of force.
- *Immediacy.* When the effects are seen within seconds to minutes—such as when a bomb explodes—the operation is probably military; if the effects take weeks or months to appear, it is more likely diplomatic or economic.
- *Directness.* If the action taken is the sole cause of the result, it is more likely to be viewed as a use of force; as the link between cause and effect attenuates, so does the military nature of the act.
- *Invasiveness.* A violated border is still an indicator of military operations; actions that are mounted from outside a target nation’s borders are probably more diplomatic or economic.
- *Measurability.* If the effect can be quantified immediately—such as photographing a “smoking hole” where the target used to be—the operation has a strong military character; the more subjective the process for evaluating the damage, the more diplomatic or economic.
- *Presumptive legitimacy.* State actors have a monopoly on the legitimate use of kinetic force, while other non-kinetic actions—attacks through or in cyberspace—are often permissible in a wider set of circumstances; actions that have not been the sole province of nation-states are less likely to be viewed as military.
- *Responsibility.* If a state takes visible responsibility for any destructive act, it is more likely to be characterized as a traditional military operation; ambiguous responsibility militates for a non-military label.

Schmitt provided two examples, each presumably premised on a state of non-hostilities existing prior to a computer network attack. In the first example, he posited computer network attacks that disable an air traffic

control (ATC) system during bad weather, resulting in the crash of an airliner and many civilian deaths. Although no kinetic force was used to destroy the airliner, CNA was the cause of the tragedy, as the airliner would have been likely to survive bad weather with a functional ATC system. The consequences are both severe and manifestly obvious, and the action (the CNA) and desired result (the airliner crash) were temporally proximate. For these reasons, this CNA can be regarded as the use of force.

In the second example, he posited a CNA on a university computer network designed to disrupt military-related research in campus laboratories. In this attack, no physical damage or suffering occurs, and the desired outcome—diminished capability on the battlefield—is both remote from the act and also depends on many other factors (e.g., the ability of researchers to regenerate data, the possible existence of other similar research efforts, and so on). In this instance, the CNA should not be regarded as the use of force.

NEW TOOLS, NEW RULES: INTERNATIONAL LAW AND INFORMATION OPERATIONS

Another more recent analysis by Duncan Hollis argued against extending traditional laws of armed conflict (LOAC) to apply to cyberattack and other information operations.² Though Hollis accepts the fundamental underlying rationale and intent of traditional LOAC (e.g., to minimize human suffering, to support reciprocity between states, to prevent morally reprehensible behavior), he argued that the interpretation of traditional LOAC vis-à-vis cyberattack suffers from two major problems.

First, Hollis argued that even in the context of state-on-state warfare, extension of the traditional LOAC suffers from serious “translation” problems about how these laws apply to cyberattack. For example, a cyberattack on a stock exchange might cause considerable economic damage but may not cause immediate death or destruction—should such an attack count as a use of force? In addition, preserving the distinction between civilian entities and valid military targets is extraordinarily difficult when cyberattack is concerned. He made the further point that Article 41 of the UN Charter defines “measures not involving the use of armed force” to include “complete or partial interruption of . . . telegraphic, radio, and other means of communication.” (Note, of course, that the UN Charter was ratified in 1945, long before the Internet and modern information

² Duncan B. Hollis, “New Tools, New Rules: International Law and Information Operations,” pp. 59-72 in *Ideas As Weapons: Influence and Perception in Modern Warfare*, G. David and T. McKeldin, eds., Potomac Books, Inc., 2009.

technologies were contemplated and before it could be imagined that the medium of an attack on a nation might well be an altogether new and different medium.)

Second, he argued that in focusing primarily on state-on-state conflict, traditional LOAC ignores many of the most important issues that arise in today's security environment—the issue of states acting against non-state actors and subnational entities. Hollis points out that the legal regimes governing such conflict are already in a state of flux (e.g., there is no doctrine comparable to the “use of force” or the self-defense provisions of the UN Charter). And when cyberattacks may be launched by non-state actors from the territories of nation-states, the relevant legal regime is even murkier.

For example, in the absence of state sponsorship, a cyberattack—even a very destructive one, conducted by a terrorist or criminal organization—does not qualify as an armed attack. A self-defense response is thus not sanctioned under the UN Charter. Even if the origin of the cyberattack can be traced to a specific state, a military or law enforcement response against an entity within that state cannot be undertaken unilaterally without violating that state's sovereignty. Only if the state in question is unable or unwilling to stop the cyberattack may the attacked state take countermeasures on its own.

Hollis concluded from his analysis that the translation difficulties and the insufficiency of traditional LOAC with respect to subnational actors call for a new legal framework for governing cyberattack and other information operations.

Appendix E

Technical Vulnerabilities Targeted by Cyber Offensive Actions

The discussion in this appendix is based largely though not entirely on an earlier National Research Council report on cybersecurity describing vulnerabilities in the information technology on which the United States relies.¹ However, there is no reason to suppose and no evidence available that suggests that other nations (or other non-national parties) are systematically better than the United States in eliminating vulnerabilities from the information technology that they use.

SOFTWARE

Software constitutes the most obvious set of vulnerabilities that an attacker might exploit. In a running operating system or application, vulnerabilities may be present as the result of faulty program design or implementation, and the exploitation of such vulnerabilities may become possible when the targeted system comes into contact with a hostile trigger (either remotely or close up). For example, a pre-implanted vulnerability in a program may be triggered at a particular time, or when a particular input is received.

When vendors find vulnerabilities, they usually issue patches to fix them. But the issuance of a patch sometimes increases the threat to those who do not install it—when a patch is widely disseminated, it also serves

¹ National Research Council, *Toward a Safer and More Secure Cyberspace*, The National Academies Press, Washington, D.C., 2007.

to notify a broad range of would-be attackers that a specific vulnerability exists. And if the patch is not installed, a broader range of attackers is likely to have knowledge of the vulnerability than if the patch had not been distributed at all. And patches are not always installed when the vendor issues them because patch installation will from time to time damage existing functionality on a system (e.g., causing a critical application to stop working until it can be made compatible with the patch to be installed).

As a rule, vulnerabilities resulting from design errors or insecure design choices are harder to fix than those resulting from implementation errors. Perhaps still more difficult are vulnerabilities introduced by unintended functionality (the euphemism for adding a function to software that helps an attacker but that is not desired by the authorized user or developer)—the classic “back-door” vulnerability.² Most system evaluation checks the extent to which a product meets the formal requirements, and not whether it does *more than* intended. Whereas vulnerabilities due to faulty design and implementation may be uncovered during the testing process or exposed during system operation and then fixed, vulnerabilities associated with unintended functionality may go undetected because the problem is tantamount to proving a negative.

Today, applications and operating systems are made up of millions of lines of code, not all of which can possibly be audited for every changed line of source code. A widely used program might have vulnerabilities deliberately introduced into it by a “rogue” programmer employed by the software vendor but planted by the attacker. (One of the most plausible vectors for the surreptitious introduction of hostile code is a third-party device driver. In some operating systems, drivers almost always require the calling system to delegate to them privileges higher than those granted

² As an example of a back door that is harmless, most versions of Microsoft Word from Word 97 to Word 2003 contain some unexpected functionality—typing “=rand()” in a Word document and then pressing the ENTER key results in three paragraphs of five repetitions of the sentence “The quick brown fox jumps over the lazy dog.” This particular back door is harmless and is even documented by Microsoft (see “How to Insert Sample Text into a Document in Word,” available at <http://support.microsoft.com/kb/212251>). Such functionality could easily not be documented, and could easily be harmful functionality as well. For example, a security interface to a computer might be designed to require the user to enter a password and to insert a physical “smart card” into a slot before granting her access. But the interface could easily be programmed to ignore the smart-card requirement when a special password is entered, and then to grant the user many more privileges than would be normal. On the other hand, the in-advance installation of a back-door vulnerability always runs a risk of premature exposure—that is, it may be discovered and fixed before the attacker can use it. Even worse from the attacker’s standpoint, it may be fixed in such a way that the attacked system appears vulnerable but is in fact not vulnerable to that particular attack. Thus, the attacker may attack and believe he was successful, even though he was not.

to ordinary users—privileges that allow the code within drivers to bypass operating system protections.) To ensure that such vulnerabilities are not introduced, vendors take many steps such as multiple code reviews during the software development process.

But source code does not always reveal the entire functionality of a system. For example, compilers are used to generate object code from source code. The compiler itself must be secure, for it could introduce object code that subversively and subtly modifies the functionality represented in the source code.³

Moreover, maliciously constructed code intentionally introduced to implant vulnerabilities in a system for later exploitation is typically more difficult to detect than are vulnerabilities that arise in the normal course of software development.⁴ Attackers highly skilled in the art of obfuscating malicious code can make finding intentionally introduced vulnerabilities a much harder problem than finding accidental flaws. Finding such vulnerabilities requires tools and skills far beyond those typically employed during system testing and evaluation aimed at discovering accidentally introduced defects. The discovery process requires detailed analysis by human experts, making it extremely expensive. Indeed, it is rarely done except for systems in which reliability and security are paramount (e.g., nuclear command and control systems).

The introduction of deliberate vulnerabilities into software is facilitated by the economic imperatives of software development and the opaqueness of the software development supply chain. Today, developing custom software for every application is impractical in terms of both cost and time. Custom software developed for a single purpose must be paid for entirely by the party for which it is developed, and thus software producers often seek to reduce costs by using commercial off-the-shelf (COTS) software and/or outsourcing their software development whenever possible (e.g., using commercial operating or database systems), even if critical systems are involved.⁵ In practice, systems are composed of components designed and implemented by many vendors. These vendors in turn often subcontract major components, and those subcontractors

³ A famous paper by Ken Thompson in 1984 described how to hide malicious binary code in a way that cannot be detected by examining the source program. See Ken L. Thompson, "Reflections on Trusting Trust," *Communications of the ACM* 27(8):761-763, August 1984.

⁴ Defense Science Board, "Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software," U.S. Department of Defense, September 2007, pp. 40-41.

⁵ Defense Science Board, "Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software," U.S. Department of Defense, September 2007, p. vi.

may in turn subcontract portions of their work. Because the spread of the Internet and high-speed communications capabilities such as broadband fiber optics worldwide has made global development of software not only possible, but also desirable for cheaply tapping the broadest range of talent,⁶ these subcontractors are often located in nations where labor is relatively inexpensive. The provenance of each component or subcomponent can only be completely known if mechanisms are in place to track each contributor, and every subcontractor represents an opportunity to introduce vulnerabilities secretly.

The use of open source software is often advocated as a solution to the security problem described above (advocates assert that the many eyes of the open source community focused on software would make it difficult or impossible to introduce deliberate flaws that will endure), and open source software is increasingly being incorporated into systems to save time and money in the development process as well. Open source software development is essentially a form of outsourced development except that the outsourcing is done on an ad hoc basis and even less may be known about the circumstances under which the code is originally produced than is the case with software produced under an outsourcing contract. Vulnerabilities could be deliberately introduced by a cyberattacker, and there is no guarantee that the open source inspection process will uncover such vulnerabilities.⁷

For example, a particular sequence of instructions and input combined with a given system state could take advantage of an obscure and poorly known characteristic of hardware functioning, which means that programmers working for an attacking government and well versed in minute behavioral details of the machine on which their code will be running could introduce functionality that would likely go undetected in any review of it.⁸

As an example of how outsourcing can be used to introduce vulnera-

⁶ Defense Science Board, "Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software," U.S. Department of Defense, September 2007, p. 10.

⁷ Empirical results appear to suggest that open source software—though available for inspection by anyone—in practice is not often audited for security. See, for example, Hal Flynn, "Why Sardonix Failed," *SecurityFocus*, February 4, 2004, available at <http://www.securityfocus.com/columnists/218>.

⁸ See, for example, Olin Sibert, Phillip A. Porras, and Robert Lindell, "An Analysis of the Intel 80x86 Security Architecture and Implementations," *IEEE Transactions on Software Engineering*, 22(5):283-293, May 1996; and Kris Kaspersky and Alice Chang, "Remote Code Execution Through Intel CPU Bugs," talk presented at Hack-In-The-Box, Dubai, United Arab Emirates, 2008, PowerPoint presentation available at <http://nchovy.kr/uploads/3/303/D2T1%20-%20Kris%20Kaspersky%20-%20Remote%20Code%20Execution%20Through%20Intel%20CPU%20Bugs.pdf>.

bilities, a financial services company reportedly outsourced its application development to a company in the Far East. The company had been certified as a CMM level-5 company, meaning that it had a well-established and documented process for developing software. However, unknown to the company, it also employed a few malicious users who inserted a back door in the application that was sent to the financial services client. The client performed only a minimal security review as part of its acceptance testing, and so the back door went undetected. The back door consisted of an undocumented URL that could be accessed remotely, through which malicious users were able to obtain customer information such as account numbers, statement balances, and other information. The back door was discovered months after deployment after the developer's clients complained about fraudulent charges.⁹

A final kind of software error is sometimes called an emergent error.¹⁰ Emergent errors can arise when correct software is used in a situation or environment for which it was not originally designed and implemented. For example, a program may work correctly in a given context and environment. However, if it is moved to a different computing environment, it may begin to work incorrectly. A software component *Z* may be certified as being secure, provided certain conditions are met (such as certain constraints on the input values being passed across its interface). It works correctly in environment *A*, which guarantees that the values passed are indeed restricted in accordance with those constraints. But if it is moved to environment *B*, which does not check the values passed to *Z*, the component may fail if values are passed that are not consistent with those constraints.

HARDWARE

Vulnerabilities can also be found in hardware, although less attention is usually paid to hardware. Hardware includes microprocessors, microcontrollers, firmware, circuit boards, power supplies, peripherals such as printers or scanners, storage devices, and communications equipment such as network cards. Tampering with such components may require physical access at some point in the hardware's life cycle, which includes access to the software and libraries of the CAD/CAM tools used to design

⁹ Ed Adams, "Biggest Information Security Mistakes That Organizations Make," Security Innovation, Inc., Wilmington, Mass., available at http://www.issa.org/Downloads/Whitepapers/Biggest-Information-Security-Mistakes_Security-Innovation.pdf.

¹⁰ Taimur Aslam, Ivan Krsul, and Eugene H. Spafford, "A Taxonomy of Security Vulnerabilities," in *Proceedings of the 19th National Information Systems Security Conference*, pp. 551-560, October 1996, available at <http://ftp.cerias.purdue.edu/pub/papers/taimur-aslam/aslam-krsul-spaf-taxonomy.pdf>.

the circuits embedded in the hardware. On the other hand, hardware is difficult to inspect, and so hardware compromises are hard to detect. Consider, for example, that peripheral devices or even other circuit cards within the main computer housing often have on-board processors and memory that can support an execution stream entirely separate from that running on a system's "main" processor.

As an experiment to demonstrate the feasibility of making malicious modifications to hardware, King et al. developed two general-purpose methods for designing malicious processors, and used these methods to implement attacks that could steal passwords, enable privilege escalation, and allow automatic logins into compromised systems.¹¹ Furthermore, the implementation of these attacks required only small amounts of modification to the baseline uncompromised processor. (For example, implementation of the login attack used only 1,341 additional logic gates, or 0.08 percent of the 1,787,958 logic gates used in the baseline; yet an attacker using this attack would gain complete and high-level access to the machine.) Embedded in larger processors involving billions of gates, the changes required would be even smaller (and thus more difficult to detect) as a percentage of the circuitry involved.

An important exception to the rule that physical access is required in order to compromise hardware is based on the fact that many systems rely on a field-upgradable read-only memory (ROM) chip to support a boot sequence, and corrupting or compromising the boot ROMs can render a system entirely non-functional (as was the case in the Chernobyl virus¹²) or only selectively non-functional. To corrupt or compromise the boot ROM that is field-upgradable, the attacker need only masquerade as a legitimate user seeking to upgrade the ROM software. Another attack on programmable read-only memory exploits the fact that the relevant chips support only a limited number of write cycles. Thus, a programmable read-only memory chip can be destroyed by an agent that repeatedly rewrites its contents a sufficient number of times. With many of today's computer system designs, corruption or destruction of a boot ROM may require at least several hours of manual repair to replace the ROM chip or some other component (such as a power supply) that may have been damaged by improper system operation. In addition, if this attack can be mounted successfully on many network routers at more or less the same time, it is likely to cause significant disruption in the overall network itself

¹¹ Samuel T. King et al., "Designing and Implementing Malicious Hardware," *Proceedings of the First USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, April 2008, available at http://www.usenix.org/event/leet08/tech/full_papers/king/king.pdf.

¹² The Chernobyl virus is further documented at http://www.cert.org/incident_notes/IN-99-03.html.

and impede network repair efforts—and so restoring the overall network to its normal operating condition will take a much longer time.

SEAMS BETWEEN HARDWARE AND SOFTWARE

Software and hardware are typically developed independently. Yet from a defensive perspective, the two are inseparable.¹³ Attacks designed to take advantage of vulnerabilities in the way software and hardware interact—almost always at some interface—may go unnoticed because testing and evaluation at the seam between them are often incidental rather than a focused activity.

COMMUNICATIONS CHANNELS

The communications channels between the system or network and the “outside” world are still another type of vulnerability. For a system to be useful it must in general communicate with the outside world, and the communications channels used can be compromised—for example, by spoofing (an adversary pretends to be the “authorized” system), by jamming (an adversary denies access to anyone else), or by eavesdropping (an adversary obtains information intended to be confidential).

One example of a communications channel cyberattack might involve seizing control of an adversary satellite by compromising its command channels. Satellites communicate with their ground stations through wireless communications, and if the command link is unencrypted or otherwise insecure, a Zendian satellite can be controlled by commands sent from the United States just as easily as by commands sent from Zendia. With access to the command link, adversary satellites can be turned off, redirected, or even directed to self-destruct by operating in unsafe modes.

CONFIGURATION

Most information technology systems—especially systems based on off-the-shelf commercial components—can be configured in different ways to support different user preferences. Configuration management—the task of ensuring that a system is configured in accordance with actual user desires—is often challenging and difficult, and errors in configuration can result in security vulnerabilities. (Many errors are the result of default configurations that turn off security functionality in order

¹³ Defense Science Board, “Report of the Defense Science Board Task Force on Mission Impact of Foreign Influence on DoD Software,” U.S. Department of Defense, September 2007, p. 4.

to ease the task of system setup. An example of such an error is a default password, such as “system” or “password,” that is widely known—such a password will remain in effect until someone chooses to change it, and such a change may never occur simply because the need to do so is overlooked.) Other configuration errors result from explicit user choices made to favor convenience—for example, a system administrator may configure a system to allow remote access through a dial-in modem attached to his desktop computer so that he can work at home, but the presence of such a feature can also be used by an attacker.

Configuration-based vulnerabilities are in some sense highly fragile, because they can be fixed on very short notice. All it takes for a configuration vulnerability to be eliminated is for the operator to choose a different configuration and implement it, which is usually a less demanding task than fixing an implementation error.

