CYBERSECURITY LAW & POLICY

INSCT
INSTITUTE FOR NATIONAL SECURITY AND COUNTERTERRORISM
SYRACUSE UNIVERSITY

Online Course FEBRUARY 13 – MARCH 30, 2012

Instructor:

Professor William C. Snyder
Visiting Assistant Professor
College of Law, Syracuse University

Reading Assignment for Module 1 – Week of 2/20 to 2/26/12

I. Introduction, Terminology, the Nature of Cyberspace and Threats

A. The Nature of Cyberspace

B. The Nature of the Threats to National Security in Cyberspace

This handout, which includes selections from:

a. Gen. Michael V. Hayden, *The Future of Things "Cyber",* Strategic Studies Quarterly (Spring 2011).
b. Joel Brenner, *America the Vulnerable: Inside the New Threat Matrix of Digital Espionage, Crime, and Warfare*, Penguin Group, Kindle Edition (2011).
c. James A. Lewis, *The Cyberwar Has Not Begun*, Center for Strategic International Studies (Mar. 2010).
d. Jonathan Zittrain, *The Future of the Internet -- And How to Stop It,* Yale University Press (2009).

The Future of Things "Cyber"

Years ago, when I was an ROTC instructor, the first unit of instruction for rising juniors dealt with communication skills. Near the beginning of the unit, I would quote Confucius to my new students: "The rectification of names is the most important business of government. If names are not correct, language will not be in accordance with the truth of things." The point had less to do with communicating than it did with thinking—thinking clearly. Clear communication begins with clear thinking. You have to be precise in your language and have the big ideas right if you are going to accomplish anything.

I am reminded of that lesson as I witness and participate in discussions about the future of things "cyber." Rarely has something been so important and so talked about with less clarity and less apparent understanding than this phenomenon. Do not get me wrong. There are genuine experts, and most of us know about patches, insider threats, worms, Trojans, WikiLeaks, and Stuxnet. But few of us (myself included) have created the broad structural framework within which to comfortably and confidently place these varied phenomena. And that matters. I have sat in *very* small group meetings in Washington, been briefed on an operational need and an operational solution, and been unable (along with my colleagues) to decide on a course of action because we lacked a clear picture of the long-term legal and policy implications of *any* decision we might make.

US Cyber Command has been in existence for more than a year, and no one familiar with the command or its mission believes our current policy, law, or doctrine is adequate to our needs or our capabilities. Most disappointingly—the doctrinal, policy, and legal dilemmas we currently face remain unresolved even though they have been around for the better part of a decade. Now is the time to think about and force some issues that have been delayed too long. This edition of *Strategic Studies Quarterly*, therefore, could not be more timely as it surfaces questions, fosters debate, and builds understanding around a host of cyber questions. The issues are nearly limitless, and many others will emerge in these pages, but let me suggest a few that frequently come to the top of my own list.

How do we deal with the unprecedented? Part of our cyber policy problem is that its newness and our familiar experience in physical space do not easily transfer to cyberspace. Casually applying well-known concepts from physical space like deterrence, where attribution is assumed, to cyberspace where attribution is frequently *the* problem, is a recipe for failure. And cyber education is difficult. In those small–group policy meetings, the solitary cyber expert often sounds like "Rain Man" to the policy wonks in the room after the third or fourth sentence. As a result, no two policymakers seemed to leave the room with the same understanding of what it was they had discussed, approved, or disapproved. So how do we create senior leaders—military and civilian who are "cyber smart enough"?

Is cyber really a domain? Like everyone else who is or has been in a US military uniform, I think of cyber as a domain. It is now enshrined in doctrine: land, sea, air, space, *cyber*. It trips off the tongue, and frankly I have found the concept liberating when I think about operationalizing this domain. But the other domains are natural, created by God, and this one is the creation of man. Man can actually change this geography, and *anything* that happens there actually creates a change in someone's *physical* space. Are these differences important enough for us to rethink our doctrine? There are those in the US government who think treating cyber as an independent domain is just a device to cleverly mask serious unanswered questions of sovereignty when conducting cyber operations. They want to be heard and satisfied before they support the full range of our cyber potential.

Privacy? When we plan for operations in a domain where adversary and friendly data coexist, we should be asking: What constitutes a twenty-first-century definition of a reasonable expectation of privacy? Google and Facebook know a lot more about most of us than we are comfortable sharing with the government. In a private-sector web culture that seems to elevate transparency to unprecedented levels, what is the appropriate role of government and the DoD? If we agree to limit government access to the web out of concerns over privacy, what degree of risk to our own security and that of the network are we prepared to accept? How do we articulate that risk to a skeptical public, and who should do it?

Do we really know the threat? Former Director of National Intelligence Mike McConnell frequently says we are already "at war" in cyberspace. Richard Clarke even titled his most recent cautionary book, *Cyber War*. Although I generally avoid the *at war* terminology, I often talk about the inherent insecurity of the web. How bad is it? And if it is really bad, with the cost of admission so low and networks so vulnerable, why have we not had a true cyber Pearl Harbor? Is this harder to do than we think? Or, are we just awaiting the inevitable? When speaking of the threat, citizens of a series of first-world nations were recently asked whom they feared most in cyberspace, and the most popular answer was not China or India or France or Israel. It was the United States. Why is that, and is it a good thing? People with money on the line in both the commercial and government sectors want clear, demonstrable answers.

What should we expect from the private sector? We all realize that most of the web things we hold dear personally and as a nation reside or travel on commercial rather than government networks. So what motivates the private sector to optimize the defense of these networks? Some have observed that the free market has failed to provide an adequate level of security for the net since the true costs of insecurity are hidden or not understood. I agree. Now what: liability statutes that create the incentives and disincentives the market seems to be lacking? Government intervention, including a broader DoD role to protect critical infrastructure beyond .mil to .gov to .com? The statutory responsibility for the latter falls to the Department of Homeland Security, but does it have the "horses" to accomplish this? Do we await catastrophe before calling for DoD intervention, or do we move preemptively?

What is classified? Let me be clear: This stuff is overprotected. It is far easier to learn about physical threats from US government agencies than to learn about cyber threats. In the popular culture, the availability of 10,000 applications for my smart phone is viewed as an unalloyed good. It is not—since each represents a potential vulnerability. But if we want to shift the popular culture, we need a broader flow of information to corporations and individuals to educate them on the threat. To do that we need to recalibrate what is truly secret. Our most pressing need is clear policy, formed by shared consensus, shaped by informed discussion, and created by a *common* body of knowledge. With no common knowledge, no meaningful discussion, and no consensus . . . the policy vacuum continues. This will not be easy, and in the wake of WikiLeaks it will require courage; but, it is essential and should itself be the subject of intense discussion. Who will step up to lead?

What constitutes the right of self defense? How much do we want to allow private entities to defend themselves outside of their own perimeters? Indeed, what should Google appropriately do *within* its own network when under attack from the Chinese state? I have compared our entry into cyberspace to mankind's last great era of discovery—European colonization of the Western Hemisphere. During that period, large private corporations like the Hudson Bay Company and the East India Tea Company acted with many of the attributes of sovereignty. What of that experience is instructive today for contemplating the appropriate roles of giants like Google and Facebook? We probably do not want to outfit twenty-first-century cyber privateers with letters of marque and reprisal, but what should be the relationship between large corporations and the government when private networks on which the government depends are under sustained attack?

Is there a role for international law? It took a decade last century for states to arrive at a new Law of the Seas Convention, and that was a domain our species had had literally millennia of experience. Then, as a powerful seafaring nation, we tilted toward maritime freedom rather than restraints. Regulating cyberspace entails even greater challenges. Indeed, as a powerful cyberfaring nation, how comfortable are we with regulation at all? After all, this domain launched by the DoD has largely been nurtured free of government regulation. Its strengths are its spontaneity, its creativity, its boundlessness. The best speech given by an American official on macro net policy was given late last year by Secretary of State Clinton when she emphasized Internet freedom, not security or control or regulation. But there are moves afoot in international bodies like the International Telecommunications Union to regulate the Internet, to give states more control over their domains, to Balkanize what up until now has been a relatively seamless global enterprise. How and when do we play?

Is cyber arms control possible? As a nation, we tend toward more freedom and less control but—given their destructiveness, their relative ease of use, and the precedent their use sets—are distributed denial-of-service attacks *ever* justified? Should we work to create a global attitude toward them comparable to the existing view toward chemical or biological weapons? Should we hold states responsible if an attack is mounted from their physical space even if there is no evidence of complicity? And, are there *any* legitimate uses for *botnets*? If not, under what authority would anyone preemptively take them down? These are questions for which no precedent in law or policy (domestic or international) currently exists. If we want to establish precedent, as opposed to likely unenforceable treaty obligations, do we emphasize dialogue with like-minded nations, international institutions . . . or multinational IT companies?

Is defense possible? At a recent conference, I was struck by a surprising question: "Would it be more effective to deal with recovery than with prevention?" In other words, is the web so skewed toward advantage for the attacker that we are reaching the point of diminishing returns for defending a network at the perimeter (or even beyond) and should now concentrate on how we respond to and recover from inevitable penetrations? This could mean more looking at *our* network for anomalous behavior than attempting to detect every incoming zero–day assault. It could mean concentrating more on what is going out rather than what is coming in. It could mean more focus on mitigating effects and operating while under attack rather than preventing attack. Mike McConnell and I met with a group of investors late last year, and we were full-throated in our warnings about the cyber threat. One participant asked the question that was clearly on everyone's mind, "How much is this going to cost me?" At the time I chalked it up to not really understanding the threat, but in retrospect our questioner may have been on to something. At what point do we shift from additional investment in defense to more investment in response and recovery?

There are more questions that could be asked, many of them as fundamental as these. Most we have not yet answered or at least have not yet *agreed* on answers, and none of them are easy. How much do we really want to empower private enterprises to defend themselves? Do we want necessarily secretive organizations like NSA or CyberCom going to the mats publicly over privacy issues? At what point does arguing for Internet security begin to legitimate China's attempts at control over Internet speech? Do we really want to get into a public debate that attempts to distinguish cyber espionage (which all countries pursue) from cyber war (something more rare and *sometimes* more destructive)? Are there any cyber capabilities, real or potential, that we are willing to give up in return for similar commitments from others?

Tough questions all—tougher (perhaps) but not unlike those our airpower ancestors faced nearly a century ago. As pioneer air warriors grappled with the unfamiliar, so must we. Until these and other questions like them are answered, we could be forced to live in the worst of all possible cyber worlds—routinely vulnerable to attack and self-restrained from bringing our own power to bear.

Gen Michael V. Hayden, USAF, Retired
Former Director, National Security Agency
Former Director, Central Intelligence Agency

America the vulnerable : inside the new threat matrix of digital espionage, crime, and warfare / Joel Brenner.

p. cm.

Includes bibliographical references and index.

ISBN : 978-1-101-54783-0

1. Computer crimes—United States—Prevention. 2. Internet in espionage—United States. 3. National security—United States.

I. Title.

HV6773.2.B74 2011

364.16’80973—dc23

2011019801

INTRODUCTION

HOW DID THE Chinese manage to remotely download up to twenty terabytes of information from the Defense Department—equal to about 20 percent of all the data in the Library of Congress? And why don’t we know exactly what they took? How did WikiLeaks get its hands on classified diplomatic cables, and why hasn’t the U.S. government been able to shut it down? How did the specifications for the avionics and armor on the president’s helicopter end up in Tehran, and what has that got to do with the theft of Supreme Court Justice Stephen Breyer’s private data from his investment adviser? The answers to these questions reveal alarming threats to our personal, corporate, and national security that come from a new type of espionage and from the sudden transparency that electronic connectivity has brought to all aspects of our lives. Your difficulties with electronic privacy, the electronic theft of America’s cutting-edge technology, and the government’s loss of state secrets are a lot more alike than you know.

I spent most of the first decade of the twenty-first century working at the heart of the U.S. government's efforts to thwart spying and terrorism against us, first as inspector general of the National Security Agency, and then as chief of counterintelligence for the director of National Intelligence. As I carried out these assignments, I saw plenty of the old-fashioned kind of espionage, but I also witnessed the dramatic rise of a new kind of spying that exploits digital technology itself, and the fact that we have all come to rely so thoroughly on that technology.

During my tenure in government I came to understand how steeply new technology has tipped the balance in favor of those—from freelance hackers to Russian mobsters to terrorists to states like China and Iran—who want to learn the secrets we keep, whether for national, corporate, or personal security. Much of my understanding arose from classified work that I cannot discuss here or anywhere. But I can share the insights I gleaned about this new form of espionage: how it works; what the biggest and most vulnerable targets are; who does it best; as well as what it means for the future of warfare, intelligence, market competition, and society at large. I also came to understand what we can—and cannot—do to counter this flood of espionage.

The truth I saw was brutal and intense: Electronic thieves are stripping us blind. I'm not just talking about the pirating of DVDs and movies in Asia or somebody ripping off your Social Security number. That's bad enough, but it's worse than that. Technologies that cost millions or billions to develop are being bled out of our corporate laboratories via the Internet; or they're slipping out after hours on thumb drives, walking onto airplanes bound for foreign ports, and reentering the country as finished products developed by foreign entrepreneurs. In effect, we're buying back our own technology. Other Western firms, meanwhile, are bleeding trade secrets, engineering designs, know-how, and other intellectual property through electronic leakage. In the public sector, sensitive diplomatic cables are suddenly splashed across the headlines worldwide. The same organizations that broadcast those cables gleefully distribute lists of critical infrastructure—airports, bridges, chemical plants—that are the most vulnerable to attack. And as I describe in the pages that follow, we're losing strategically sensitive data about aircraft and ship design, radars, and other defense technology, as well as information about auto manufacturing, engineering designs, and other commercial innovations. This theft contributes to the tidal flow of capital from West to East that threatens our prosperity, and it could in wartime cost many American lives.

[redacted]

[redacted]

[redacted]

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

The Cyber War Has Not Begun

James Andrew Lewis
Center for Strategic and International Studies

March 2010

Expanded attention to cybersecurity is a good thing, but it seems that it is difficult to discuss this topic without exaggeration. We are not in a 'cyber war'. War is the use of military force to attack another nation and damage or destroy its capability and will to resist. Cyber war would involve an effort by another nation or a politically motivated group to use cyber attacks to attain political ends. No nation has launched a cyber attack or cyber war against the United States.

Indeed, it would be a bold nation that would do so. A deliberate attack on the United States could trigger a violent if not devastating response. No nation would be foolish enough to send a missile, aircraft or commando team to attack critical infrastructure in this country. The same logic applies to cyber attack. Foreign leaders will not lightly begin a war with the United States and the risk of cyber war is too high for frivolous or spontaneous engagement.

[redacted]

[redacted]

[redacted]

[redacted]

- [redacted]

THE FUTURE OF THE INTERNET

AND HOW TO STOP IT

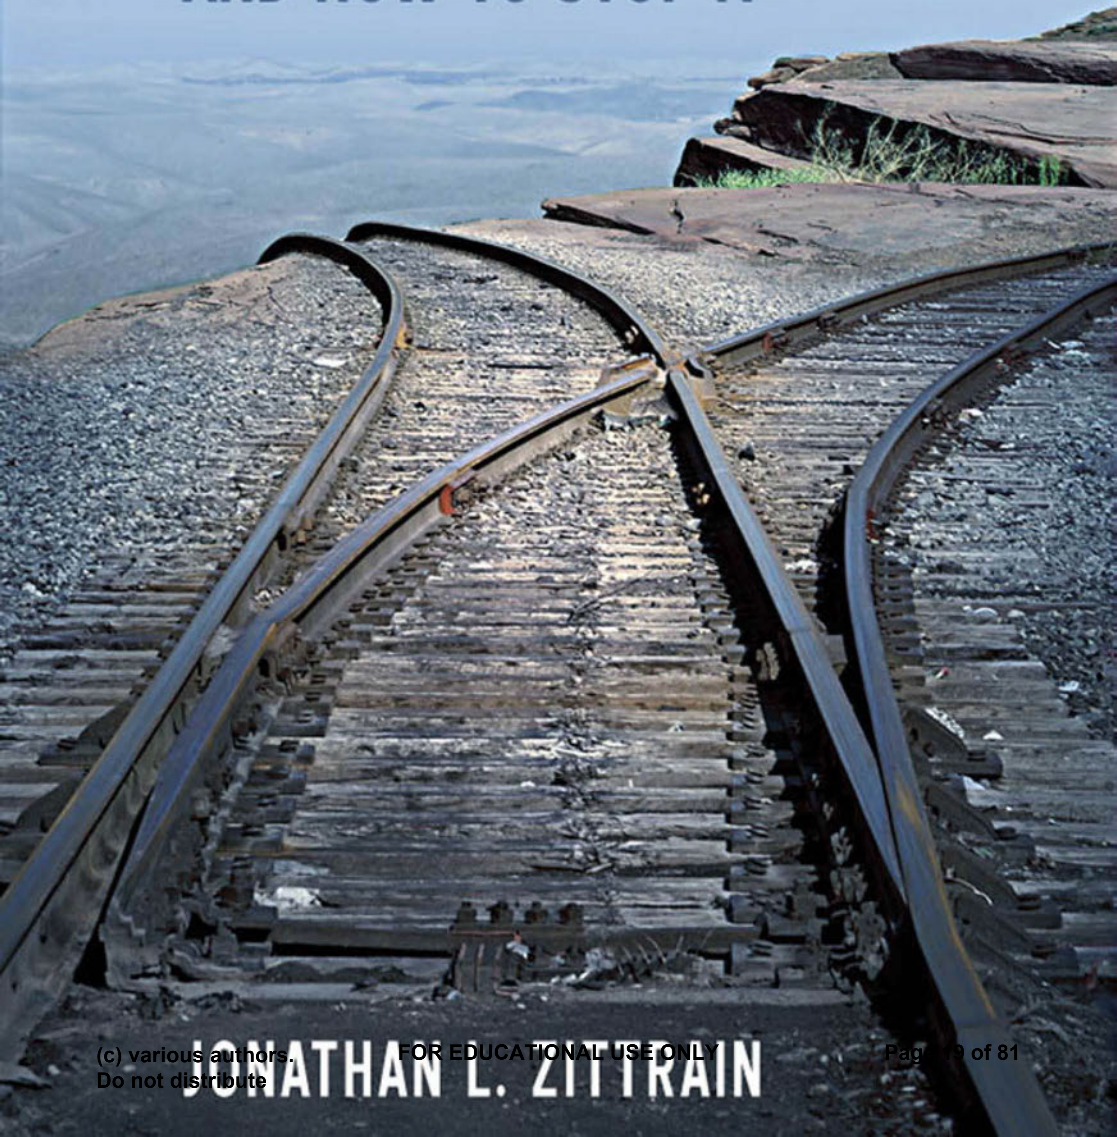

JONATHAN L. ZITTRAIN

A Caravan book. For more information, visit www.caravanbooks.org.

The cover was designed by Ivo van der Ent, based on his winning entry of an open competition at www.worth1000.com.

Set in Adobe Garamond type by The Composing Room of Michigan, Inc.

Printed in the United States of America by R. R. Donnelley, Harrisonburg, Virginia.

Library of Congress Control Number: 2008942463
ISBN 978-0-300-15124-4 (pbk. : alk. paper)

A catalogue record for this book is available from the British Library.

10 9 8 7 6 5 4 3 2 1

Preface to the Paperback Edition

The venerable Warner Brothers antagonist Wile E. Coyote famously demonstrates a law of cartoon physics. He runs off a cliff, unaware of its ledge, and continues forward without falling. The Coyote defies gravity until he looks down and sees there's nothing under him. His mental gears turn as he contemplates his predicament. Then: splat.

Both the Internet and the PC are on a similar trajectory. They were designed by people who shared the same love of amateur tinkering as the enterprising Coyote. Both platforms were released unfinished, relying on their users to figure out what to do with them—and to deal with problems as they arose. This kind of openness isn't found in our cars, fridges, or TiVos. Compared to the rest of the technologies we use each day, it's completely anomalous, even absurd.

This openness, described and praised in this book in more detail as "generativity," allowed the Internet and PC to emerge from the realms of researchers and hobbyists and surprisingly win out over far more carefully planned and funded platforms. (They were certainly more successful than any of the Coyote's many projects.)

Today the very popularity and use of the Internet and PC are sorely testing that generativity. We wouldn't want our cars, fridges, or TiVos to be altered by unknown outsiders at the touch of a button—and yet this remains the prevailing way that we load new software on our PCs. More and more often that software is rogue—harvesting computing cycles from a PC in order to attack others, stealing personal information, or simply frying the PC. Soon, either abruptly or in slow motion: splat.

The first reaction to abuses of openness is to try to lock things down. One model for lockdown can be drawn from our familiar appliances, which are sealed when they leave the factory. No one but a true geek could hack a car or a fridge—or would want to—and we've seen glimpses of that model in communications platforms like iPods, most video game consoles, e-book readers like the Amazon Kindle, and cable company set-top boxes. Such lockdown was the direction a visionary Steve Jobs—the guy who gave us the first open PC, the Apple II—first took with the iPhone, with which he bet the future of Apple.

Of course, the Internet or PC would have to be in bad shape for us to abandon them for such totally closed platforms; there are too many pluses to being able to do things that platform manufacturers don't want or haven't thought of. But there's another model for lockdown that's much more subtle, and that takes, well, a book to unpack. This new model exploits near-ubiquitous network connectivity to let vendors change and monitor their technologies long after they've left the factory—or to let them bring us, the users, to them, as more and more of our activities shift away from our own devices and into the Internet's "cloud."

These technologies can let geeky outsiders build upon them just as they could with PCs, but in a highly controlled and contingent way. This is iPhone 2.0: an iPod on steroids, with a thriving market for software written by outsiders that must be approved by and funneled through Apple. It's also Web 2.0 software-as-service ventures like the Facebook platform and Google Apps, where an application popular one day can be banished the next.

This model is likely the future of computing and networking, and it is no minor tweak. It's a wholesale revision to the Internet and PC environment we've experienced for the past thirty years. The serendipity of outside tinkering that has marked that generative era gave us the Web, instant messaging, peer-to-peer networking, Skype, Wikipedia—all ideas out of left field. Now it is disappearing, leaving a handful of new gatekeepers in place, with us and them prisoner to their limited business plans and to regulators who fear things that are new and disruptive. We are at risk of embracing this new model, thinking it the best of both worlds—security and whimsy—when it may be the worst. Even fully grasping how untenable our old models have become, consolidation and lockdown need not be the only alternative. We can stop that future.

Introduction

On January 9, 2007, Steve Jobs introduced the iPhone to an eager audience crammed into San Francisco's Moscone Center.[1] A beautiful and brilliantly engineered device, the iPhone blended three products into one: an iPod, with the highest-quality screen Apple had ever produced; a phone, with cleverly integrated functionality, such as voicemail that came wrapped as separately accessible messages; and a device to access the Internet, with a smart and elegant browser, and with built-in map, weather, stock, and e-mail capabilities. It was a technical and design triumph for Jobs, bringing the company into a market with an extraordinary potential for growth, and pushing the industry to a new level of competition in ways to connect us to each other and to the Web.

This was not the first time Steve Jobs had launched a revolution. Thirty years earlier, at the First West Coast Computer Faire in nearly the same spot, the twenty-one-year-old Jobs, wearing his first suit, exhibited the Apple II personal computer to great buzz amidst "10,000 walking, talking computer freaks."[2] The Apple II was a machine for hobbyists who did not want to fuss with soldering irons: all the ingredients for a functioning PC were provided in a convenient molded plastic case. It looked clunky, yet it could be at home on someone's desk. Instead of puzzling over bits of hardware or typing up punch cards to feed into someone else's mainframe, Apple owners faced only the hurdle of a cryptic blinking cursor in the upper left corner of the screen: the PC awaited instructions. But the hurdle was not high. Some owners were inspired to program the machines themselves, but true beginners simply could load up software written and then shared or sold by their more skilled or inspired counterparts. The Apple II was a blank slate, a bold departure from previous technology that had been developed and marketed to perform specific tasks from the first day of its sale to the last day of its use.

The Apple II quickly became popular. And when programmer and entrepreneur Dan Bricklin introduced the first killer application for the Apple II in 1979—VisiCalc, the world's first spreadsheet program—sales of the ungainly but very cool machine took off dramatically.[3] An Apple running VisiCalc helped to convince a skeptical world that there was a place for the PC at everyone's desk and hence a market to build many, and to build them very fast.

Though these two inventions—iPhone and Apple II—were launched by the same man, the revolutions that they inaugurated are radically different. For the technology that each inaugurated is radically different. The Apple II was quintessentially *generative* technology. It was a platform. It invited people to tinker with it. Hobbyists wrote programs. Businesses began to plan on selling software. Jobs (and Apple) had no clue how the machine would be used. They had their hunches, but, fortunately for them, nothing constrained the PC to the hunches of the founders. Apple did not even know that VisiCalc was on the market when it noticed sales of the Apple II skyrocketing. The Apple II was designed for surprises—some very good (VisiCalc), and some not so good (the inevitable and frequent computer crashes).

The iPhone is the opposite. It is sterile. Rather than a platform that invites innovation, the iPhone comes preprogrammed. You are not allowed to add programs to the all-in-one device that Steve Jobs sells you. Its functionality is locked in, though Apple can change it through remote updates. Indeed, to those who managed to tinker with the code to enable the iPhone to support more or different applications,[4] Apple threatened (and then delivered on the threat) to transform the iPhone into an iBrick.[5] The machine was not to be *generative* beyond the innovations that Apple (and its exclusive carrier, AT&T) wanted. Whereas the world would innovate for the Apple II, only Apple would innovate for the iPhone. (A promised software development kit may allow others to program the iPhone with Apple's permission.)

Jobs was not shy about these restrictions baked into the iPhone. As he said at its launch:

> We define everything that is on the phone. . . . You don't want your phone to be like a PC. The last thing you want is to have loaded three apps on your phone and then you go to make a call and it doesn't work anymore. These are more like iPods than they are like computers.[6]

No doubt, for a significant number of us, Jobs was exactly right. For in the thirty years between the first flashing cursor on the Apple II and the gorgeous iconized touch menu of the iPhone, we have grown weary not with the unexpected cool stuff that the generative PC had produced, but instead with the unexpected very uncool stuff that came along with it. Viruses, spam, identity theft, crashes: all of these were the consequences of a certain freedom built into the generative PC. As these problems grow worse, for many the promise of security is enough reason to give up that freedom.

* * *

In the arc from the Apple II to the iPhone, we learn something important about where the Internet has been, and something more important about where it is going. The PC revolution was launched with PCs that invited innovation by others. So too with the Internet. Both were generative: they were designed to accept any contribution that followed a basic set of rules (either coded for a particular operating system, or respecting the protocols of the Internet). Both overwhelmed their respective proprietary, non-generative competitors, such as the makers of stand-alone word processors and proprietary online services like CompuServe and AOL. But the future unfolding right now is very different from this past. The future is not one of generative PCs attached to a generative network. It is instead one of sterile *appliances* tethered to a network of control.

These appliances take the innovations already created by Internet users and package them neatly and compellingly, which is good—but only if the Internet and PC can remain sufficiently central in the digital ecosystem to compete with locked-down appliances and facilitate the next round of innovations. The balance between the two spheres is precarious, and it is slipping toward the safer appliance. For example, Microsoft's Xbox 360 video game console is a powerful computer, but, unlike Microsoft's Windows operating system for PCs, it does not allow just anyone to write software that can run on it. Bill Gates sees the Xbox as at the center of the future digital ecosystem, rather than at its periphery: "It is a general purpose computer. . . . [W]e wouldn't have done it if it was just a gaming device. We wouldn't have gotten into the category at all. It was about strategically being in the living room. . . . [T]his is not some big secret. Sony says the same things."[7]

It is not easy to imagine the PC going extinct, and taking with it the possibility of allowing outside code to run—code that is the original source of so much of what we find useful about the Internet. But along with the rise of information appliances that package those useful activities without readily allowing new ones, there is the increasing lockdown of the PC itself. PCs may not be competing with information appliances so much as they are becoming them. The trend is starting in schools, libraries, cyber cafés, and offices, where the users of PCs are not their owners. The owners' interests in maintaining stable computing environments are naturally aligned with technologies that tame the wildness of the Internet and PC, at the expense of valuable activities their users might otherwise discover.

The need for stability is growing. Today's viruses and spyware are not merely annoyances to be ignored as one might tune out loud conversations at nearby tables in a restaurant. They will not be fixed by some new round of patches to bug-filled PC operating systems, or by abandoning now-ubiquitous Windows for Mac. Rather, they pose a fundamental dilemma: as long as people control the code that runs on their machines, they can make mistakes and be tricked into running dangerous code. As more people use PCs and make them more accessible to the outside world through broadband, the value of corrupting these users' decisions is increasing. That value is derived from stealing people's attention, PC processing cycles, network bandwidth, or online preferences. And the fact that a Web page can be and often is rendered on the fly by drawing upon hundreds of different sources scattered across the Net—a page may pull in content from its owner, advertisements from a syndicate, and links from various other feeds—means that bad code can infect huge swaths of the Web in a heartbeat.

If security problems worsen and fear spreads, rank-and-file users will not be far behind in preferring some form of lockdown—and regulators will speed the process along. In turn, that lockdown opens the door to new forms of regulatory surveillance and control. We have some hints of what that can look like. Enterprising law enforcement officers have been able to eavesdrop on occupants of motor vehicles equipped with the latest travel assistance systems by producing secret warrants and flicking a distant switch. They can turn a standard mobile phone into a roving microphone—whether or not it is being used for a call. As these opportunities arise in places under the rule of law—where some might welcome them—they also arise within technology-embracing authoritarian states, because the technology is exported.

A lockdown on PCs and a corresponding rise of tethered appliances will eliminate what today we take for granted: a world where mainstream technology can be influenced, even revolutionized, out of left field. Stopping this future depends on some wisely developed and implemented locks, along with new technologies and a community ethos that secures the keys to those locks among groups with shared norms and a sense of public purpose, rather than in the hands of a single gatekeeping entity, whether public or private.

The iPhone is a product of both fashion and fear. It boasts an undeniably attractive aesthetic, and it bottles some of the best innovations from the PC and Internet in a stable, controlled form. The PC and Internet were the engines of those innovations, and if they can be saved, they will offer more. As time passes, the brand names on each side will change. But the core battle will remain. It will be fought through information appliances and Web 2.0 platforms like today's Facebook apps and Google Maps mash-ups. These are not just products but also services, watched and updated according to the constant dictates of their makers and those who can pressure them.

In this book I take up the question of what is likely to come next and what we should do about it.

I

The Rise and Stall of the Generative Net

Today's Internet is not the only way to build a network. In the 1990s, the Internet passed unnoticed in mainstream circles while networks were deployed by competing proprietary barons such as AOL, CompuServe, and Prodigy. The technorati placed bets on which baron would prevail over the others, apparently imagining that the proprietary networks would develop in the same way that the separate phone networks—at one time requiring differently colored phones on each person's desk—had converged to just one lucky provider.[1] All those bets lost. The proprietary networks went extinct, despite having accumulated millions of subscribers. They were crushed by a network built by government researchers and computer scientists who had no CEO, no master business plan, no paying subscribers, no investment in content, and no financial interest in accumulating subscribers.

The framers of the Internet did not design their network with visions of mainstream dominance. Instead, the very unexpectedness of its success was a critical ingredient. The Internet was able to develop quietly and organically for years before it became widely known, remaining outside the notice of those who would have insisted on more cautious strictures had they only suspected how ubiquitous it would become.

This first part of the book traces the battle between the centralized proprietary networks and the Internet, and a corresponding fight between specialized information appliances like smart typewriters and the general-purpose PC, highlighting the qualities that allowed the Internet and PC to win.

Today, the same qualities that led to their successes are causing the Internet and the PC to falter. As ubiquitous as Internet technologies are today, the pieces are in place for a wholesale shift away from the original chaotic design that has given rise to the modern information revolution. This counterrevolution would push mainstream users away from a *generative* Internet that fosters innovation and disruption, to an *appliancized* network that incorporates some of the most powerful features of today's Internet while greatly limiting its innovative capacity—and, for better or worse, heightening its regulability. A seductive and more powerful generation of proprietary networks and information appliances is waiting for round two. If the problems associated with the Internet and PC are not addressed, a set of blunt solutions will likely be applied to solve the problems at the expense of much of what we love about today's information ecosystem. Understanding its history sheds light on different possible futures and helps us to recognize and avoid what might otherwise be very tempting dead ends.

One vital lesson from the past is that the endpoint matters. Too often, a discussion of the Internet and its future stops just short of its endpoints, focusing only on the literal network itself: how many people are connected, whether and how it is filtered, and how fast it carries data.[2] These are important questions, but they risk obscuring the reality that people's experiences with the Internet are shaped at least as much by the devices they use to access it.

As Internet-aware devices proliferate, questions posed about network regulation must also be applied to the endpoints—which, until recently, have been so open and so nonconstricting as to be nearly unnoticeable, and therefore absent from most debates about Internet policy. Yet increasingly the box has come to matter.

History shows that the box had competitors—and today they are back. The early models of commercial (as compared to academic) computing assumed that the vendor of the machinery would provide most or all of its programming. The PC of the 1980s—the parent of today's PC—diverged from these models, but the result was by no means a foregone conclusion. Internet users are again embracing a range of "tethered appliances," reflecting a resurgence of the initial model of bundled hardware and software that is created and controlled by one company. This will affect how readily behavior on the Internet can be regulated, which in turn will determine the extent that regulators and commercial incumbents can constrain amateur innovation, which has been responsible for much of what we now consider precious about the Internet.[3]

The Internet also had competitors—and they are back. Compared to the Internet, early online information services were built around very different technical and business models. Their designs were much easier to secure against illegal behavior and security threats; the cost was that innovation became much more difficult. The Internet outpaced these services by assuming that every user was contributing a goodwill subsidy: people would not behave destructively even when there were no easy ways to monitor or stop them.

The Internet's tradeoff of more flexibility for less security worked: most imaginable risks failed to materialize—for example, people did not routinely spy on one another's communications, even though it was eminently possible, and for years there were no spam and no viruses. By observing at which point these tradeoffs were made, we will see that the current portfolio of tradeoffs is no longer optimal, and that some of the natural adjustments in that balance, while predictable, are also undesirable.

The fundamental challenges for those who have built and maintained the Internet are to acknowledge crucial deficiencies in a network-and-endpoint structure that has otherwise served so well for so long, to understand our alternatives as the status quo evaporates, and to devise ways to push the system toward a future that addresses the very real problems that are forcing change, while preserving the elements we hold most dear.

[redacted]

the generative PC.[111] That is the generative dilemma. The next chapter explains more systematically the benefits of generativity, and Chapter Five explores what the digital ecosystem will look like should our devices become more thoroughly appliancized. The vision is not a pleasant one, even though it may come about naturally through market demand. The key to avoiding such a future is to give that market a reason not to abandon or lock down the PCs that have served it so well—also giving most governments reason to refrain from major intervention into Internet architecture. The solutions to the generative dilemma will rest on social and legal innovation as much as on technical innovation, and the best guideposts can be found in other generative successes in those arenas. Those successes have faced similar challenges resulting from too much openness, and many have overcome them without abandoning generativity through solutions that inventively combine technical and social elements.