CERIAS Tech Report 2004-06

TRUSTED COMPUTING: THE DEBATE OVER MAKING CYBERSPACE SAFE FOR COMMERCE

James C. Hinde

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907-2086

TRUSTED COMPUTING:
THE DEBATE OVER
MAKING CYBERSPACE
SAFE FOR COMMERCE

A Thesis

Submitted to the Faculty

of

Purdue University

by

James C. Hinde, Jr.

In Partial Fulfillment of the

Requirements for the Degree

of

Master of Science

May 2004

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Hinde, James C., Jr. M.S., Purdue University, May, 2004. Trusted Computing: The Debate Over Making Cyberspace Safe For Commerce. Major Professor: Toby J. Arquette.

"Trusted computing" refers to a pair of architectural standards proposed in 1999 and 2002 by the Trusted Computing Platform Alliance and Microsoft Corporation, respectively, for a new generation of computers and operating systems that would incorporate into their operation the use of cryptographic techniques to validate the software running on those systems. These proposals, especially the one from Microsoft under the code name "Palladium," were met with angry protests from many quarters. Critics charged that trusted computing was an attempt to force digital rights management upon the public and would open the door to unprecedented intrusions upon the freedom and privacy of the users of personal computers.

This thesis examines the controversy over trusted computing. It identifies the prominent features of the historical context in which the ideas of digital rights management and trusted computing arose in the 1990s. It documents the positions of both sides of the debate that followed the news of Palladium. Having found that debate to be polarized and unproductive, the thesis examines its ideological underpinnings and attempts to point the way to a more honest, imaginative and useful debate.

INTRODUCTION

The cover article of the July 1, 2002, issue of *Newsweek* introduced the public to a newly-announced project at Microsoft Corporation, which author Steven Levy (2002) described as "Microsoft's hyperambitious long-range plan to literally change the architecture of PCs in order to address the concerns of security, privacy and intellectual property." Code-named Palladium after the statue of Athena that according to legend guarded the ancient city of Troy, this project proposed to implement and commercialize for the first time an approach to computer security known as *trusted computing*, a term whose ambiguity would soon become emblematic of a profound controversy that has yet to be resolved. Two years earlier, in October 1999, Microsoft, Compaq, IBM, Intel and Hewlett-Packard had founded the Trusted Computing Platform Alliance, or TCPA, a non-profit standards body whose mission was to develop a set of specifications for implementing trusted computing on a wide variety of devices that store and transmit digital information, including personal data assistants and cellular telephones as well as computers. The TCPA published the first version of its specification in July 2001. Palladium was expected to incorporate some of the TCPA specification's standards into its own design, but also to add features that were not included in the TCPA specification.

Because the debate over trusted computing has unfolded as a historical process, the first three chapters of this thesis, which will review the relevant literature, will be organized for the most part chronologically. Chapter 1 will look at documents that define

the TCPA's conception of what trusted computing is and describe the scientific, legal and economic context in which it developed up to the time of Microsoft's announcement of Palladium. Chapter 2 will look at Microsoft's own early statements about Palladium and the first wave of critical reaction to both Palladium and the TCPA during the remaining months of 2002. Chapter 3 will examine the industry's reaction to these criticisms, as articulated in documents produced by Microsoft and the Trusted Computing Group (or TCG, the successor organization to the TCPA) in 2003. It will be seen that the debate of trusted computing to date has been highly polarized and characterized more by posturing and name-calling than by dialog. The final three chapters of the thesis will attempt to clear the air around the debate by identifying some of the ideological undercurrents that surround and obscure it. Chapter 4 will introduce a small number of theoretical works in the fields of law, sociology and philosophy that offer insights into the broader social, political and economic issues that are at play in the debate. Chapter 5 will apply those theoretical tools to the task of illuminating the trusted computing debate, and the concluding chapter will recommend some measures that should be taken to allow the debate to move forward in a productive and honest manner.

CHAPTER 1: THE HISTORICAL CONTEXT

Past trusted systems

The concept of trusted systems did not originate with the TCPA. Computer scientists in academia and government had grappled with the problem of trust in the security of computer systems for many years prior to 1999. The Department of Defense (1985) published its definition of trusted systems in the Trusted Computer System Evaluation Criteria (TCSEC). Commonly known as the Orange Book, the TCSEC defined a detailed set of design and implementation criteria that were to be used in evaluating the degree to which a system could be expected to resist unauthorized attempts to read, alter or block access to the information it contained, along with a set of seven ratings classes to indicate the results of the evaluation of specific systems. Windows NT, for example, achieved a TCSEC rating of C2, the third-lowest of the seven classes. While its details are beyond the scope of this review, the TCSEC was noteworthy for the present discussion because it defined the "trustworthiness" of a computer system as something that could be ascertained through a formal evaluation of the engineering methods used in its design and development. Because networked computer systems had not yet come into widespread use in the early 1980s, the TCSEC was developed in the mainframe-centered context of the 1970s, where systems were typically monolithic. In addition to running on a single computer, a system of that era presented little ambiguity as to the identity and needs of its owner. The relatively straightforward evaluation

approach that this permitted gave way to exponential increases in complexity, difficulty and expense when distributed systems appeared in the early 1990s.

Landwehr (1993) spoke to this concern in a paper that proposed a streamlined, simplified and integrated approach to security evaluation. This addressed the existence of distributed applications in which hardware and software components from multiple sources were combined in a modular fashion to perform a business function, but Landwehr did not question the TCSEC's underlying assumption that a system's trustworthiness was conceptually uncomplicated.

In the same year Denning (1993) proposed a radical redefinition of the concept of trust in computer systems, based on the phenomenon of trust in everyday life:

> It is an assessment that a person, organization or object can be counted on to perform according to a given set of standards in some domain of action. As an assessment, it is a declaration made by an observer rather than an inherent property of the person, organization or object observed (p. 37).

Denning further insisted that in the case of a computer system the assessment of trust that matters is the one made by the users of the system, and that the mechanism by which users declare their assessments of a system is their informed decision in an efficient market to purchase or not to purchase that system. This paper anticipated later developments in the industry both in its emphasis on market forces and in placing the spotlight on the users as the determiners of whether, to what degree and for what purposes a system is trusted.

Jøsang (1996) expanded on Denning's concept of trust as an assessment by attempting to catalog a variety of types of trust relationships that might exist around computer systems. Among Jøsang's insights was the concept of "origin diversity of trust" (p. 126), which recognized that different trusting entities could trust a single target in different ways. Because this paper was primarily theoretical, Jøsang did not provide any real-world examples of origin diversity.

Perhaps the most succinct and definitive statement about the problem of placing trust in a computer system was made by Thompson (1984) in an article with the provocative title "Reflections on Trusting Trust". After leading the reader through a demonstration of how a language compiler could be surreptitiously modified to place a Trojan horse inside an application program compiled from legitimate source code, Thompson concludes: "The moral is obvious. You can't trust code that you did not totally create yourself" (p. 763). Any technical certification of the trustworthiness of any piece of code can, at some level of processing or at some stage in the development process, be circumvented, falsified or otherwise defeated. Trusting the code inevitably boils down to trusting all the people who participated in its creation.

Copyright legislation

While computer scientists were grappling with the difficulty of evaluating the trustworthiness of systems in the 1990s, the entertainment industry was growing concerned about the trustworthiness of the general public. Three technological trends

were converging to produce a situation that for the industry was simultaneously an unprecedented opportunity and an unprecedented danger. Analog media products—records, audio cassettes, film photography, video cassettes, and books—were increasingly being superceded by digital products that could be consumed or produced using devices that could easily be connected to or built into inexpensive personal computers. Tens of millions of people were connecting their home computers to the Internet, and high bandwidth connections were becoming both widely available and easily affordable by many consumers. Finally, software and services were being made accessible over the Internet which allowed users to easily compress, exchange and share music and video files. These developments were creating a situation that was both a blessing and a curse for the entertainment industry. They were benefiting the industry by creating an infrastructure for direct and efficient distribution in real time of digital entertainment products. This infrastructure promised to realize the dream of video on demand, in which any consumer could at any hour of the day connect to a studio's Web site and, for a fee, download and watch any movie in that studio's catalog. However, the same trends were also making it possible for ordinary people using ordinary household equipment to have the capability to make and distribute over the Internet high quality copies of copyrighted digital products that were already on the market. Commercialization of video on demand would be impossible until a way could be found to prevent the downloaded content from being pirated.

One avenue of effort the entertainment industry followed to address the piracy problem was to lobby Congress to pass new legislation that would expand the rights of

intellectual property owners. The three most significant pieces of copyright legislation passed in the late 1990s were the No Electronic Theft (NET) Act (1997), the Sonny Bono Copyright Term Extension Act (CTEA) (1998), and the Digital Millennium Copyright Act (DMCA) (1998).

The NET Act closed a loophole that had been revealed in an unsuccessful prosecution in 1994. David LaMacchia, a student at MIT who had freely distributed unlicensed copies of copyrighted software from a bulletin board service that he operated, had been charged with conspiracy to commit wire fraud, but the case against him was dismissed because it could not be demonstrated that he had profited financially from that activity. The NET Act removed the requirement of profit from the definition of criminal copyright infringement. The witnesses who testified before the U. S. House of Representatives Subcommittee on Courts and Intellectual Property (1997) in its hearing on NET included representatives of the Motion Picture Association of America, the Software Publishers Association, the Recording Industry Association of America, Microsoft Corporation and Adobe Systems, Incorporated.

The CTEA extended the terms of most existing copyrights by twenty years. Ostensibly this was done to bring United States copyright law into conformity with European law. But because one of the beneficiaries of the extension was the Walt Disney Co., whose copyright on the first film in which Mickey Mouse appeared would otherwise have expired in 2003, the CTEA was widely perceived as having by passed at Disney's behest. The Chicago Tribune (Disney Lobbying for Copyright Extension, 1998) and the Washington Post (McAllister, 1998) separately reported at the time that Disney CEO

Michael Eisner had personally lobbied both Senate Majority Leader Trent Lott and House Speaker Newt Gingrich in favor of the CTEA, and that Disney had made significant campaign contributions to most of the sponsors of the bill in both houses of Congress.

Samuelson (2000) reports that the DMCA was in most respects a direct outgrowth of a white paper created by the Clinton Administration's Working Group on Intellectual Property Rights under the title "Intellectual Property and the National Information Infrastructure" (Working Group on Intellectual Property Rights, 1995). This paper voiced the concern that tightening copyright law might not in itself be sufficient to prevent widespread infringement in the age of the Internet. Because this concern lies at the heart of the idea of using technology to carry out enforcement measures that have traditionally been the province of the legal system, the pertinent passage from the Working Group paper is worth reproducing here:

> The ease of infringement and the difficulty of detection and enforcement will cause copyright owners to look to technology, as well as the law, for protection of their works. However, it is clear that technology can be used to defeat any protection that technology may provide. The Working Group finds that legal protection alone will not be adequate to provide incentive to authors to create and to disseminate works to the public. Similarly, technological protection likely will not be effective unless the law also provides some protection for the technological processes and systems used to prevent or restrict unauthorized uses of copyrighted works.

The primary purpose of the DMCA was to criminalize the circumvention of technological mechanisms whose purpose was to protect copyrighted digital material from unauthorized copying. The law also made it a crime to "manufacture, import, offer to the public, provide, or otherwise traffic in any technology, product, service, device, component, or part thereof" that is designed for the purpose of accomplishing such a circumvention or has limited usefulness for any other purpose. Its passage in 1998 was made possible by the inclusion of a "safe harbor" provision that exempted online service providers from liability for infringements perpetrated by their customers.

The NET Act, the CTEA and the DMCA were most significant pieces of intellectual property legislation passed in the late 1990s, but they represent only a small portion of Congressional interest and activity in that field during those years. The United States Copyright Office (2004) lists on its Web site a total of 18 copyright-related bills introduced in the 105th Congress, 19 in the 106th and 22 in the 107th.

Defenders of the public domain

A number of critics voiced concern over the changes to copyright law in the 1990s. They included cyberspace visionaries, constitutional scholars, and some writers who were both at once. What united these critics was the belief that progress in the information age depended upon the free exchange of information and was threatened by a hasty, ill-considered and one-sided attempt on the part of Congress to promote and protect the parochial interests of the software, music and film industries.

Decades before the invention of personal computers or even timesharing mainframe systems, visionaries dreamed of a world in which computer technology would make the collective knowledge of mankind available to individuals. In a remarkably prescient article Vannevar Bush (1945) imagined a future device for which he invented the name *memex*, and which functioned in much the same way that the personal computer would half a century later. Although Bush imagined that people would find a wide variety of uses for the memex, he saw it as above all a tool for extending scientific progress. Science, he wrote, "has provided a record of ideas and has enabled man to manipulate and to make extracts from that record so that knowledge evolves and endures through the life of a race rather than that of an individual" (p. 101). The accumulation of knowledge in the common record was, Bush thought, threatened by the size to which that record had grown, creating a logjam of ideas that hindered efficient research and forced scientists into artificially narrow specialties. The memex would break that logjam and "yet allow him [i.e. "man"] truly to encompass the great record and to grow in the wisdom of race experience" (p. 108).

As years passed later in the century, the specific problems which computers were expected to solve would vary, but the dream of collaborative information sharing remained central to the idea of progress. Licklider (1968) defined the interactive computer as primarily a communication device, and Nelson (1974) envisaged a comprehensive system in which all human knowledge would be indexed and made available online. During the 1970s the first personal computers and electronic bulletin board services were developed in a countercultural atmosphere documented by Roszak

(1993) as “electronic populism”. The Community Memory project, established in Berkeley and San Francisco in 1973 (Felsenstein, 1993), was a deliberate attempt to create an “electronic agora” in which ordinary people could freely share the information stored in a timeshared computer. From their beginnings in bulletin board services through the advent of the Internet, networks of personal computers connected through the public telecommunications infrastructure were frequently described as an “electronic frontier” (Rheingold, 1993; Barlow, 1994) that resembled the American West of the nineteenth century. According to this view, the rules, values, power structures and property interests of the real-life world—the same “straight” world the counterculture had attempted to escape from a decade earlier—were as irrelevant to cyberspace as the tired and corrupt ways of the East had been to the Old West (Barlow, 1996). In its most extreme form, this kind of thinking took the form of the Hacker ethic, which flatly denied the legitimacy of any property rights whatsoever to data stored in systems connected to the Internet.

Opposition to the expansion of intellectual property rights over digital information was also forthcoming from less radical quarters during the years leading up to Microsoft’s Palladium announcement. Several legal and technology scholars (Boyle, 2002; Benkler, 2001; Lessig, 2002a; Davis, 2001) expressed concern that the new copyright regime was defeating the purpose of the copyright clause in Article 1, Section 8 of the United States Constitution, which authorizes Congress to “promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.” These writers argued that

the effectiveness of copyright for promoting progress depends upon a careful balance being struck between the rights of past inventors to be compensated for their work and the ability of present and future inventors to build upon the stock of existing inventions to create new ones. This balance, they further argued, was being tilted in favor of past inventors, with the result that new invention would be inhibited by the obligation of researchers to identify and comply with an endless and bewildering array of ownership claims to existing ideas.

This argument was taken to the United States Supreme Court in the Eldred v. Ashcroft (2003) case, in which the plaintiff claimed that the CTEA was unconstitutional on the grounds that Congress was evading the constitutional requirement that copyrights be granted only for a limited time. Although the Court upheld the constitutionality of the CTEA, the majority opinion was less an endorsement of the law's wisdom than a refusal to second-guess Congress on its merits. The dissenting opinions by Justices Stevens and Breyer make it clear that for them at least, the value of the public domain for future innovation was the central issue in the case. Stevens wrote:

> By failing to protect the public interest in free access to the products of inventive and artistic genius—indeed, by virtually ignoring the central purpose of the Copyright/Patent Clause—the Court has quitclaimed to Congress its principal responsibility in this area of the law (p.60).

In addition to their concern for the public domain, the legal critics of increased copyright protection also shared with Bush and the radical hackers the opinion that advances in information technology presented society with an opportunity for achievement that

required for its realization that old ways of doing things be abandoned. Benkler (2001) cast the argument in a vocabulary of biological evolution, describing the technological innovations of the 1990s as "shocks to the economic and technological ecosystem" that had caused "new species of information production" to appear, in the forms of open source software developers and peer-to-peer file sharing networks. Against these newcomers, the previously dominant species—Hollywood, the music industry and publishers of proprietary software—were ill adapted to compete in the new environment. The efforts of these institutions to shore up copyright protection, Benkler argued, were harmful not only to the public interest but to the private interests of those same institutions, who would be better off devoting their energy to adapting to the new world than to doomed efforts to preserve the old one. Davis (2001) suggested the development of new business models as the best solution to the problem of unauthorized copying. The possibilities he suggested were all designed to make legitimately purchased digital products sufficiently more attractive than illicit copies that customers would prefer purchasing to copying.

Technological defenses against infringement

By the end of the 1990s Hollywood's reluctance to make its products available for video on demand distribution across the Internet was becoming a business problem for the companies that wanted to provide the computer hardware and software that would make that kind of distribution possible. Personal computers had progressed to the point

where further increases in speed and storage capacity could not be exploited by traditional applications in ways that would persuade consumers to buy replacements for current models. Home users in particular had little use for PCs that would recalculate spreadsheets, spell-check documents and display Web pages twice as fast as the ones they already had. The ability to download movies from the Internet and watch them on the computer, on the other hand, would go a long way toward justifying the expense of an upgrade. Although the information technology industry supported the recent copyright legislation—Microsoft, for example, having sent a representative to testify in favor of passing the NET Act—that industry differed from Hollywood in that it had both the ability and the inclination to treat piracy as a technical problem.

In 2001, Microsoft secured a patent (United States Patent Office, 2001) for a "digital rights management operating system." According to the copyright notice in the patent application, Microsoft wrote the application in 1998. The patent application's abstract, reproduced below in its entirety, reprises the terminology of trust from the computer science literature reviewed earlier in this chapter.

> A digital rights management operating system protects rights-managed data, such as downloaded content, from access by untrusted programs while the data is loaded into memory or on a page file as a result of the execution of a trusted application that accesses the memory. To protect the rights-managed data resident in memory, the digital rights management operating system refuses to load an untrusted program into memory while the trusted application is executing or removes the data from memory before loading the untrusted program. If the

> untrusted program executes at the operating system level, such as a debugger, the digital rights management operating system renounces a trusted identity created for it by the computer processor when the computer was booted. To protect the rights-managed data on the page file, the digital rights management operating system prohibits raw access to the page file, or erases the data from the page file before allowing such access. Alternatively, the digital rights management operating system can encrypt the rights-managed data prior to writing it to the page file. The digital rights management operating system also limits the functions the user can perform on the rights-managed data and the trusted application, and can provide a trusted clock used in place of the standard computer clock.

In these six sentences the words *trusted* or *untrusted* appear nine times. The proposed operating system is designed to distinguish between trusted and untrusted programs and to make it impossible, when a trusted program and an untrusted program are both running on the computer at the same time, for the latter to access "rights-managed data" that may be in use by the former. A trusted program, as defined later in the text of the patent application, is one that has been "authenticated as respecting digital rights," and the only rights under discussion are "the content providers' rights." These rights are not defined or limited by copyright law, but instead have been spelled out contractually in a license agreement between the content provider and the computer's owner. It can be presumed that when the owner originally installed the trusted program on that system, he or she was

required to agree to the license agreement's terms before the installation could be completed.

The trust relationship that lies at the heart of the digital rights management operating system is one that was not included in Jøsang's (1996) catalog. Here, the owner of the digital information to be used on the computer trusts some of the software on that computer—the operating system and the application program that will use the information in question—to enforce the terms of the license agreement. Such terms might include not allowing a copy to be made of the information, only allowing the information to be viewed a certain number of times, erasing the information at the end of a certain length of time, or anything else. The untrusted parties in this arrangement include not only any other programs that might be running on that computer but also the person at the keyboard. This is stated explicitly in the patent application: "In a very real sense, the legitimate user of a computer can be an adversary of the data or content provider. 'Digital rights management' is therefore fast becoming a central requirement if online commerce is to continue its rapid growth."

The trust relationship between the software and the content provider is supported by an attestation process that runs whenever the system is booted and in which a cryptographic digest of the operating system is created and checked to verify that the operating system has not been altered (in any way that might defeat its digital rights management capability) since its original installation, followed by a similar attestation of each trusted application program that is run during the session. Because these attestation processes can detect alteration of the trusted software caused by a virus infection or the

introduction of a Trojan horse program, it is possible to argue that they provide some assurance to the computer's user that the computer is "secure," but that assurance is limited by the extent to which the user's definition of security coincides with that of the content provider. Later, Microsoft would use that argument in its efforts to win consumer acceptance of the Palladium proposal.

The TCPA

Within months of its formation in 1999, the TCPA posted two short white papers on its web site that purport to describe the organization's vision, its guiding principles and the general outlines of the technical specification that it would publish at a later date. The first paper (TCPA, 2000a) states that the TCPA's goal is "to build a solid foundation for trust in the PC over time," and that "the specification for the trusted PC platform should focus on two areas—ensuring privacy and enhancing security" (p. 1). The second paper (TCPA, 2000b) defines trusted computing in terms of authenticity, integrity and privacy and states that "A central objective of the TCPA specification is to protect privacy by maintaining owner control over critical data" (p. 4).

The operation of a computer that complies with the TCPA specification, as described in both of these papers, is similar to the operation of the digital rights management operating system for which Microsoft had filed a patent application in 1998 (United States Patent Office, 2001), in that it incorporates an attestation mechanism that a connected computer could use to verify, before transmitting any data across the

connection, that the software running on the subject computer is "trusted." Its most significant departure from Microsoft's earlier proposal is that the TCPA attestation would be performed through the use of a tamper-proof hardware module that would be built into every TCPA-compliant computer.

Each of these papers presents a small number of business scenarios designed to demonstrate the benefits of using TCPA-compliant systems. Only one scenario is present in both papers: the scenario in which a PC provides attestation to a server that its software is trusted before the server permits it to download proprietary data.

The picture of trusted computing painted by the two TCPA papers is less one-sided in favor of the content provider than the one in Microsoft's patent application. These papers include language about the need for systems to be trusted by users as well as by servers, and there is some discussion of ways that the specification could preserve personal privacy by permitting a client computer to attest itself to a server without revealing its owner's identity. There is a kind of specious see-no-evil neutrality in these papers, which seeks to avoid responsibility for any abuses that the system might facilitate. For example, the description of remote attestation in the first paper includes the statement: "It is important to note that the TCPA process does not make value judgments regarding the integrity metrics" (p. 5). In other words, if a content provider refuses to trust any media player other than the one sold by its business partner, the unapproved player is just as unacceptable as a Trojan horse, and that is not the TCPA's fault. In the end, however, it is clear that the principal beneficiaries of trusted computing will be large organizations and digital content providers.

The CBDTPA

Three months after Microsoft received its patent for a digital rights management operating system, a bill was introduced in the United States Senate that would require digital rights management capabilities to be built into any device sold in the United States that is capable of storing or transmitting copyrighted digital information (CBDTPA, 2001). Sponsored by Senator Ernest Hollings (D-SC), the Consumer Broadband and Digital Television Promotion Act, or CBDTPA, would have, if enacted, made the sale of any digital device that fails to include and utilize security technologies that "provide for secure technical means of implementing directions of copyright owners for copyrighted works" a federal crime punishable by 5 years in prison and a $500,000 fine. In the statement he made upon introducing the bill (Hollings, 2001), Senator Hollings stated its rationale with a clarity that would become politically untenable within a year:

> The fact is that most Americans are averse to paying $50 a month for faster access to email, or $2000 for a fancy HDTV set that plays analog movies. But if more high-quality content were available, consumer interest would likely increase.
>
> By unleashing an avalanche of digital content on broadband Internet connections as well as over the digital broadcast airwaves, we can change this dynamic and give consumers a reason to buy new consumer electronics and information technology products. To do so requires the development of a secure, protected environment to foster the widespread dissemination of digital content in these exciting new mediums.

> Although it is technologically feasible to provide such a protected environment, the solution has not been forthcoming through voluntary private sector negotiations involving the industries with stakes in this matter. This is not to say, however, that those industries do not recognize the tremendous economic potential to be derived from a proliferation of top notch digital content to consumers in the home. The movie studios, and the rest of the copyright industries, for example, are tremendously excited about the possibility of providing their products to consumers over the Internet and the digital airwaves, provided they can be assured that those products' copyrights are not infringed in the process.
>
> Although marketplace negotiations have not provided such an assurance, a solution is at hand. Leaders in the consumer electronics, information technology, and content industries are some of America's best and brightest. They can solve this problem.

The private sector, Senator Hollings believed, needed a nudge to make them give consumers a reason to buy broadband Internet connections and digital television sets. The CBDTPA failed to make it out of the Judiciary Committee (McCullagh, 2002) and therefore never gave the information technology industry the intended nudge, but it did attract the attention of software professionals, security experts and civil libertarians. Eight days after the bill's introduction, the co-chairs of the U.S. Public Policy Committee of the Association for Computing Machinery (USACM) sent a letter to Senator Hollings (Simons and Spafford, 2002) in which they said that the provisions of the CBDTPA

would be both ineffective against piracy and detrimental to a wide variety of legitimate uses of digital computing. The letter also pointed out that entertainment was "only one, relatively minor use (compared to all uses) of networks and computing technology," and urged the Senator to take a broader view. As will be seen later in this review, opponents of digital rights management were not the only people who were making this observation.

In a less polite manner characteristic of many people who were outraged by the CBDTPA, security expert Richard Forno (2002) posted a broadside on the Internet in which he described Senator Hollings as the "Senator from Disney" and a member of the "American Techniban." The proponents of the CBDTPA, wrote Forno, in terms reminiscent of Benkler's (2001) evolutionary metaphor, wanted "to effect electronic martial law on all information resources and implement draconian measures on today's information society for no other reason than to satisfy the profiteering desires of the entertainment moguls desperately trying to save their crumbling Industrial Age business models."

Information security concerns

The widespread adoption of the Internet as a communications medium in the 1990s placed large numbers of people in the uncomfortable new position of having to worry about their computers being attacked by viruses, worms, Trojan horses and other malicious software. Because many of these attacks were made possible by defects in

commercial software products, the vendors of those products were frequently blamed by the press and the public for failure to build safe products.

No software vendor came under more fire for security vulnerabilities than Microsoft. A historical review of computer viruses and attacks conducted by the *Washington Post* (Krebs, 2003) listed seven major incidents that occurred during the years 1998 through 2001. Each of these had been a major news story at the time of its appearance, and five of them—Melissa, I Love You, Anna Kournikova, Code Red and Nimda—involved the exploitation of vulnerabilities in Microsoft email, word processing or Web server products. For the same four years, the ICAT Metabase (NIST, 2004) shows that as the total number of reported software vulnerabilities for all products increased each year, Microsoft's share of that count never fell below 10 per cent.

The spreading perception that its products were responsible for security breaches was more than a public relations problem for Microsoft. In January of 2002, the National Academy of Science published a report (NRC, 2002) that recommended that software vendors be held legally liable for damages that could be traced to defects in their products. Microsoft was under increasing pressure to convince the world that it could be trusted to clean up its products' security weaknesses without the need for government intervention.

	1998	1999	2000	2001
Total reported vulnerabilities	245	862	990	1506
Reported vulnerabilities in Microsoft products	25	157	139	157
Percentage of total	10%	18%	14%	10%

Table 1. Vulnerabilities reported by ICAT, 1998 through 2001

Summary

By the summer of 2002 the debate over trusted systems and digital property rights had been joined from many directions. Computer scientists were wrestling with the question of what it means to trust a computer. Hollywood was exploring every avenue it could think of to prevent digital piracy. The computer industry was hoping to see digital content delivery drive the demand for a new generation of hardware and software. Hackers and civil libertarians were rallying to defend the electronic frontier and the public domain. People in the press, the information security community and the government were complaining about the security weaknesses in commercial software. If one were to describe all these ideas and points of view as a set of ingredients that were being churned at a slow speed setting in a giant blender, it would be fair to say that Microsoft chairman Bill Gates was about to switch that blender to a higher speed.

CHAPTER 2: ANNOUNCEMENTS AND CRITICAL REACTIONS

In the *Newsweek* article that announced Palladium to the world (Levy, 2002), Bill Gates is quoted as saying, “It’s a funny thing. We came at this thinking about music, but e-mail and documents were far more interesting domains” (p. 49). As the USACM had urged Senator Hollings to do, Microsoft had recognized that entertainment was only a small part of the picture. The digital rights management technology that had been developed for the purpose of defeating piracy of music and videos could be applied to any type of digital data, and the range of “rights” that trusted applications could enforce was limited only by programmers’ imaginations. One example mentioned in the *Newsweek* article was an e-mail client program that would obey instructions associated with each incoming message that could prohibit the recipient from printing the message, forwarding it to unapproved third parties, copying its contents to another document, or even looking at it again after a specified length of time.

Microsoft was not able to control the timing or the manner in which Palladium was announced to the world. As one reporter told the story,

> Microsoft tried to keep a lid on the story for as long as possible. But after finding out that Levy was going to print something, the company invited him to Redmond for two days to hear the whole story. Even then, Microsoft didn't expect the story to run so soon. When it discovered that Levy's story was about to hit the

> streets, Microsoft barely had time to warn those of us who were maintaining our silence that the secret was almost out of the bag.
>
> I'm telling you all this because Microsoft would have been better off staying silent on this one. The reports that are surfacing are going to raise many more questions than Microsoft has answers for. (Coursey, 2002)

Microsoft attempted to position Palladium to the public as a solution to computer users' concerns about security. Within weeks of the *Newsweek* article, Bill Gates (2002) published an "Executive E-mail" on Microsoft's Web site, under the title "Trustworthy Computing." This five-page document listed a number of recent corporate initiatives at Microsoft which, Gates claimed, demonstrated that the company, after having been criticized for years for marketing products that were riddled with security vulnerabilities, had made security its highest priority. The list included a short paragraph about Palladium:

> We are working on a new hardware/software architecture for the Windows PC platform, code-named "Palladium," which will significantly enhance users' system integrity, privacy and data security. This new technology, which will be included in a future version of Windows, will enable applications and application components to run in a protected memory space that is highly resistant to tampering and interference. This will greatly reduce the risk of viruses, other attacks, or attempts to acquire personal information or digital property with malicious or illegal intent.

In its ambiguity, characteristic of arguments for trusted computing, about who is trusting the system to protect whose information from access by whom, this statement did little to stem the torrent of adverse comment that had already broken loose upon the publication of the *Newsweek* article. Headlines that appeared on Web news outlets within days included: "Who trusts Microsoft's Palladium? Not me" (Loney, 2002) on June 27, "Is Microsoft's Palladium a Trojan Horse?" (Morrissey, 2002), on June 28, "Why we can't trust Microsoft's 'trustworthy' OS" (Coursey, 2002) on July 2, "Control Your Identity or Microsoft and Intel Will" on July 9, and "Can we trust Microsoft's Palladium?" (Manjoo, 2002) on July 11.

The criticism aroused by the news of Palladium owed its vehemence to two primary factors. First, the news that Palladium would use digital rights management techniques on email messages and document files made it a concern for many people for whom the use of DRM on entertainment material was, if not something they approved of, at most a minor irritant. Palladium threatened to make it impossible for office workers everywhere to send email to their friends from work. Second, the announcement that Palladium would be incorporated into an upcoming release of Windows added enough of a sense of impending reality to what had been until then a theoretical debate, to focus the minds of commentators on the connections between elements that until then had seemed only loosely related to one another. Trusted computing, tightened copyright laws, the criminalization of reverse-engineering anything that could be claimed to be a copy protection mechanism, not to mention a suddenly heightened awareness of how little we really know about what goes on inside our computers—all these crystallized for many

critics into an Orwellian nightmare. In *Nineteen Eighty-Four*, Winston Smith's job at the Ministry of Truth had been to rewrite history to suit the government by replacing offending paper documents with new versions and sending the originals down a Memory Hole. With Palladium's help, it suddenly seemed, he would have been able with a click of a mouse to enlist the cooperation of every computer in the world to do the job electronically.

Before reviewing the specific criticisms that were leveled against Palladium in the wake of its announcement, it will be useful to examine briefly its features and the relationship between Palladium and the TCPA specification, as Microsoft described them in a white paper (Microsoft, 2002a) and a frequently-asked-questions (FAQ) list (Microsoft, 2002b) posted in August 2002. The FAQ identified four features as central to Palladium:

> **Q: What is the "Palladium" initiative, anyway?**
>
> A: The "Palladium" code name refers to both hardware and software changes. Specifically, it refers to a new set of features in the Microsoft® Windows® operating system that, when combined with new hardware and software, provide additional security services to PCs. There are four categories of these features:
>
> - **Curtained memory.** The ability to wall off and hide pages of main memory so that each "Palladium" application can be assured that it is not modified or observed by any other application or even the operating system

- **Attestation.** The ability for a piece of code to digitally sign or otherwise attest to a piece of data and further assure the signature recipient that the data was constructed by an unforgeable, cryptographically identified software stack
- **Sealed storage.** The ability to securely store information so that a "Palladium" application or module can mandate that the information be accessible only to itself or to a set of other trusted components that can be identified in a cryptographically secure manner
- **Secure input and output.** A secure path from the keyboard and mouse to "Palladium" applications, and a secure path from "Palladium" applications to a region of the screen

When running, "Palladium" provides a parallel execution environment to the "traditional" Windows kernel- and user-mode stacks; "Palladium" runs alongside the OS, not underneath it.

The goal with "Palladium" is to help protect software from software; that is, to provide a set of features and services that a software application can use to defend against malicious software also running on the machine (viruses running in the main operating system, keyboard sniffers, frame grabbers, etc). "Palladium" is not designed to provide defenses against hardware-based attacks that originate from someone in control of the local machine.

The FAQ states further that the sealed storage and attestation features are what Palladium and the TCPA specification have in common. (Later on, Microsoft narrowed the difference between the two specifications by stating (Microsoft, 2003b) that the hardware

component of the NGSCB—Palladium's second incarnation—would adhere to the TCPA specification.) The definition of attestation in the passage cited above is obscure to a non-technical reader. A more function definition is given in the white paper: "Attestation is a mechanism that allows the user to reveal selected characteristics of the operating environment to external requestors. For example, attestation can be used to verify that the computer is running a valid version of 'Palladium.'" Similarly, attestation could be used to verify to a content provider that the computer is running an approved media player.

It should be noted that attestation as proposed by Palladium and the TCPA is closely related but not identical to integrity checking, a widely used information security mechanism which also seeks to verify that unauthorized modifications have not been made to software stored on a computer's hard disk. The difference between the two is that the power to decide which software is authorized is left by integrity checking in the hands of the computer's owner, whereas under Palladium or TCPA it belongs to the holder of digital rights over the data that the user is attempting to access.

Because the curtained memory and secure input/output features of Palladium were neither controversial nor shared with the TCPA specification, and because Microsoft was a charter member of the TCPA, the critics of trusted computing generally did not distinguish between the Palladium and TCPA. The criticisms of trusted computing that arose in 2002 fall roughly into five categories, which are covered in the sections below.

Customer lock-in

Anderson (2002c), Arbaugh (2002), and Stallman (2002) argued that trusted computing would make it impossible for open source operating systems or application programs to run as trusted software, because trusted computing requires a trusted component to be digitally signed and certified by an authority as trustworthy. That certification process is predicated upon the assumption that any alteration of a program renders it untrustworthy, whether it is the result of a software attack or of a legitimate improvement made by the computer's owner. Certifying any open source program would effectively freeze its ongoing development, because any changes made to it beyond that point would make it untrusted until a new certification could be performed, requiring a process that no individual programmer could afford to undertake.

Anderson (2002c), Green (2002b), and Stallman (2002) point out that at the application level, trusted computing would enable software vendors to prevent competitors from supporting their proprietary file formats. A trusted version of Microsoft Word running under Palladium, for example, could create documents that would trust only another copy of Word to read them. Over time, a user or an organization would accumulate an inventory of documents which would all become unreadable if the application used to create them were replaced by a competitor's product. To make matters worse, trusted computing would enable software vendors to rent licenses to customers instead of selling them, leading to the automatic loss of access to all documents, no matter where they might be located, that were created by a copy of an application whose license rental fee is not paid on time. Any attempt to recover the

content of a file that has been rendered unreadable in this way could be construed as a violation of the DMCA.

Censorship

Censorship can operate proactively, by making it impossible for an author to distribute forbidden information, or retroactively, by retrieving and destroying all existing copies of a banned work after it has been published and distributed. Stallman (2002) provides an example of how trusted computing, which he pointedly calls "treacherous computing," could support the former in a whistle-blowing situation:

> Imagine if you get an email from your boss stating a policy that is illegal or morally outrageous, such as to shred your company's audit documents, or to allow a dangerous threat to your country to move forward unchecked. Today you can send this to a reporter and expose the activity. With treacherous computing, the reporter won't be able to read the document; her computer will refuse to obey her. Treacherous computing becomes a paradise for corruption.

The threat of retroactive censorship is perhaps the most frightening scenario that critics associated with trusted computing. Stallman (2002) provides a succinct description of how the remote attestation feature of trusted computing could permit the retroactive erasure of documents:

> Programs that use treacherous computing will continually download new authorization rules through the Internet, and impose those rules automatically on

> your work. If Microsoft, or the US government, does not like what you said in a document you wrote, they could post new instructions telling all computers to refuse to let anyone read that document. Each computer would obey when it downloads the new instructions. Your writing would be subject to 1984-style retroactive erasure. You might be unable to read it yourself.

This scenario, complete with its reference to Orwell, is based upon a claim made by Lucky Green (2002b) in a presentation entitled *TCPA: the mother(board) of all Big Brothers*, that he gave at the DefCon X conference in Las Vegas in August, 2002. Speaking to the slide reproduced in Figure 1 below, Green said that trusted computing would support a "document revocation list" (DRL) for every trusted application program capable of producing document files, that each local copy of each application would query its DRL over the Internet as part of the remote attestation process, and that no trusted application would open a document file that was on its DRL. Because retroactive

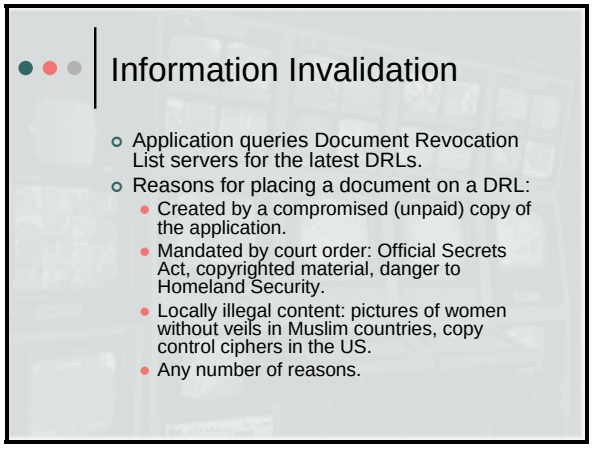

Figure 1. Slide 29 from Lucky Green's DefCon X presentation

censorship was the most sinister behavior that anyone accused trusted computing of making possible, because the claim was vehemently denied by Microsoft, and because the debate over DRLs illustrates the atmosphere in which the Palladium announcement was received by the public, that debate is worth looking at in some detail

The first mention of DRLs in connection with trusted computing is in an email from Lucky Green to the Cryptography mailing list and Ross Anderson (Green, 2002a) on June 26, 2002. Referring to the statement in the *Newsweek* article that Palladium would support the creation of "Word documents that could only be read in the next week" (Levy, 2002, p. 49), Green begins his argument by pointing out that such a capability requires that the trusted application program have access to a secure clock that cannot be reset by a user attempting to fool the system into thinking an expired document is still unexpired. Arguing that the secure clock would need to reside on an Internet server, he concludes that trusted computing applications will need to regularly access the Internet in order to operate. Because what transpires during that Internet access is entirely under the control of the application software vendor, Green's argument continues, it is reasonable to suppose that it could include the downloading of blacklists of pirated copies of application software. Because every document created by a trusted application can be traced to the individual copy of the program used to create it, the blacklist could be used to invalidate every document ever created by any unlicensed copy of the trusted application. Finally, the Internet access could also include the downloading of a list of individual documents that for any reason have been deemed unacceptable and not to be read. Green concludes:

> All that is required to perform such an administrative invalidation of a document is either a sample copy of the document from which one can obtain its globally unique ID, the serial number of the application that created the document, or the public key of the person who licensed the application.

Although Green's June 26 email does not include the term "document revocation list", a post by Anderson (2002a) two days later credits Green for having invented the term: "I discussed this with Lucky Green yesterday and he came up with the delightful term 'document revocation list'." Apparently Anderson was intrigued as well as delighted by Green's idea, because he claims in a second post (Anderson, 2002b) one day after the first to have confirmed that Microsoft had plans for a DRL up its corporate sleeve:

> I've been in Stanford the last few days for a conference and a number of people have been giving me more information on TCPA/ Palladium. It is clear that it was from the start a DRM project (Bill admits this) and there also appear to have been plans from the start to include the 'document revocation list' idea.

Anderson subsequently incorporated the idea of "remote censorship" into the first version of his "Trusted Computing Frequently Asked Questions" Web site (Anderson, 2002c), which rapidly came to be regarded by many people as an authoritative source of information about trusted computing. (As a rough measure of this, the number of Web sites that, according to a Google search on March 22, 2004, contain links to the current version of Anderson's FAQ site is 2,300.)

Green's and Anderson's assertions about retroactive censorship drew denials from Microsoft and angry reaction from at least one quarter. Microsoft's FAQ page

(Microsoft, 2002b) flatly denied that retroactive disabling of unlicensed software was part of Palladium:

> Q: Some people have claimed that “Palladium” will enable Microsoft or other parties to detect and remotely delete unlicensed software from my PC. Is this true?
>
> A: No. As stated above, the function of “Palladium” is to make digitally signed statements about code identity and hide secrets from other “Palladium” applications and regular Windows kernel- and user-mode spaces. “Palladium” doesn't have any features that make it easier for an application to detect or delete files.

A poster to the Cryptography mailing list who used the name “AARG!Anonymous” (2002) took Green and Anderson to task for creating a distorted picture of both Palladium and the TCPA. He called Green’s DefCon presentation “a tissue of lies and fabrications and unfounded sensationalism” and accused one of his own critics of supporting the “ridiculous claims in Ross Anderson's FAQ” out of unwillingness “to publicly take a position in opposition to such a famous and respected figure.” Later in the same thread, another poster (“Bear”, 2002) accused “AARG” of being the dishonest one:

> The spec is designed to be hard to read and M. AARG, the one who has been talking about the advantages of the proposal, has been (hmmm) either terribly naive or deliberately misleading. As people actually work through the spec and find the things s/he's been claiming aren't there or couldn't be done with it, the

> odds of his/her being a mere paid shill increase and his/her credibility decreases in direct proportion.

The week following Green's DefCon presentation, Green appeared along with Peter Biddle, the Microsoft manager widely regarded as the architect of Palladium, in a panel discussion about Palladium at the 11th USENIX Security Symposium in San Francisco (USENIX, 2002). The third member of the panel was Seth Schoen, the Staff Technologist at the Electronic Frontier Foundation. Schoen, who had attended Green's DefCon presentation and had also been a member of a delegation from the EFF who were invited by Microsoft to meet with the Palladium team, described both of these experiences in his personal Web log (Schoen, 2002a and 2002b) and wrote a white paper about trusted computing for the EFF (Schoen, 2003). He turned out to be a rare voice of moderation in the midst of a highly polarized debate, and his testimony is valuable both for explaining the positions of the partisans to neutral outsiders and for his own measured critique of trusted computing. The position he takes with respect to the DRL debate is that both sides are right:

> Lucky predicts with some confidence that Palladium will be used to do things like the DRL. But it's also the case that the DRL is not 'part of Palladium'. As far as I can tell, it's something which application vendors would be able to implement under Palladium" (Schoen, 2002b).

Schoen proceeds to describe in detail a protocol that a trusted application program and a DRL server could use under Palladium to implement document revocation, and he acknowledges that Green is not unreasonable to suppose that application vendors might

be pressured by their large institutional customers to add revocation services to their products.

The controversy over the DRL did not make Anderson back down from his claim that trusted computing would lead to retroactive censorship. It did, however, lead him to state his position more carefully and at greater length in the updated version of his trusted computing FAQ list (Anderson, 2003) that he posted in August, 2003, where he suggested that the first cases of retroactive censorship might be ordered by courts in child pornography cases.

Closely related to the problem of censorship is the potential for criminals to use the trusted computing architecture to keep evidence of illegal activity beyond the reach of law enforcement agencies. A trader in child pornography, for example, could encrypt incriminating files so that only his customers could open them. If so, the authorities would be prevented not only from using those files as evidence but in many cases from being able to learn that the crime is being committed in the first place.

Prevention of first sale and fair use

United States copyright law includes two statutory limitations on the rights of copyright holders, the first sale doctrine and the fair use doctrine. The first sale doctrine permits the owner of a legally purchased copy of a copyrighted work, such as a book, to sell or lend that copy to a third party without securing the copyright holder's permission to do so. The fair use doctrine permits people to make and distribute, subject to standards

that have been worked out in the courts, partial copies of copyrighted works for educational use or for critical evaluation.

Schneier (2002) points out that the use by Palladium (which he abbreviates by using the element's chemical symbol, Pd) of cryptographic keys could effectively prohibit a user from selling his or her copy of a trusted application program.

> Like books and furniture and clothing, the person who currently buys new software can resell it when he's done with it. People have a right to do this -- it's called the "First Sale Doctrine" in the United States -- but the software industry has long claimed that software is not sold, but licensed, and cannot be transferred. When someone sells a Pd-equipped computer, he is likely to clear his keys so that his identity can't be used or files can't be read. This will also serve to erase all the software he purchased. The end result might be that people won't be able to resell software, even if they wanted to.

For copyrighted content such as music or movies, trusted computing's digital rights management features could support a transfer of the license to a file from one user to another, by requiring the seller's copy to be erased as part of the transaction. Whether content providers would permit such an arrangement is an open question.

Fair use presents a difficult if not impossible problem for trusted computing, because to permit fair use without opening the door to piracy a trusted application would need to be able to know the user's intentions. Lessig (2002b) expressed doubt that Palladium applications would protect fair use rights. Olsen (2003) reported that in the academic community, where fair use currently permits professors to copy short reading

assignments for their students, many professors fear that Palladium may curtail those rights, and she also quotes an official of the Business Software Alliance as complaining that colleges “aren't sending the message as aggressively as we would like” about copyright infringement on campuses.

Invasion of privacy

As noted earlier (TCPA, 2000a), the TCPA claimed from the beginning of their effort that ensuring the privacy of users was a primary goal of their effort. The measures outlined in the TCPA specification for preserving privacy all revolve around the idea of allowing the client computer to authenticate itself to a trusted third party, who would then give it a credential that it could use to persuade a server that it meets the latter’s criteria for trustworthiness without needing to reveal the user’s identity to the server. Arbaugh (2002) argues that this arrangement is inherently flawed because there is no guarantee that the trusted third party will not misuse the information entrusted to it by users. Anderson (2002c) argues that specification’s claim to champion privacy is a red herring, given that in real life “almost all privacy violations result from the abuse of authorised access, often obtained by coercing consent.”

In a later development, an agency of the European Union raised, in a report issued in January 2004 (EU, 2004), concerns about the use of a trusted third party to safeguard users’ personal identifiable information in the latest specification from the TCPA’s successor, the Trusted Computing Group. That report and the TCG’s response a few weeks later (TCG, 2004) indicate that the TCG has been willing to amend its

specification in response to criticism, at least when the criticism comes from entities that possess regulatory power.

Failure to protect against viruses

The attestation and curtained memory mechanisms of both Palladium and the TCPA specification, it is claimed, would make it difficult for virus-infected software to damage a system. Attestation would cause a program that had been altered by a virus infection to fail to be certified as trusted, with the result that it would not be permitted to access data that required a trusted application. Curtained memory would ensure that the memory within which a trusted application was running would be inaccessible to any untrusted program running on the same computer that would seek to infect the trusted application with a virus. Anderson (2002c) and Lasser (2002) point out, however, that the most common virus attacks in recent years do not attempt to alter executable programs but rather to exploit the scripting capabilities of popular desktop applications such as Microsoft Outlook and Microsoft Word. Because a virus of this type does not require the host program to be attacked by an external process, curtained memory would not prevent a trusted application from reading a document (assuming that all the licensing requirements were met) that contained a script directing the application to misbehave in some way. From the end user's point of view, remote attestation itself could be considered a form of macro virus.

Summary

The first wave of reaction to the Palladium announcement was concerned much more with what trusted computing would permit powerful interests to do than with what it might do by itself. For the critics, what was at stake was the libertarian dream of personal empowerment, autonomy and freedom that had been a driving force in the computer revolution since its inception. In 1984, Apple had introduced the Macintosh computer with a television commercial during the Super Bowl in which a young woman challenged an Orwellian totalitarian regime by hurling a hammer into Big Brother's telescreen. For the next two decades, the hippies and the yuppies had found common cause in making interactive computing an integral part of middle-class life in large parts of the world. The alliance had developed numerous signs of strain by 2002. Microsoft in particular, the young company that had beaten IBM and ushered in the new age of personal computing, had come under suspicion from many quarters, of having grown up to be a greedy and arrogant world power. But prior to the Palladium announcement, nothing Microsoft had done—not its anticompetitive business practices, not its overwhelming market share, not the astronomical wealth of its executives—had been perceived as a direct attack on individual users. Coming from a company that had embodied the libertarian dream for many people, Microsoft's embrace of trusted computing and digital rights management was felt as more than an attack. It was felt as a betrayal.

CHAPTER 3: THE INDUSTRY'S RESPONSE

Neither Microsoft nor the TCPA answered the critics of trusted computing in a substantive way. Instead of arguing that trusted computing would not permit the restrictions on user autonomy that the critics said it would, or arguing that such restrictions were a social price worth paying for the benefits of more efficient electronic commerce, both Microsoft and the TCPA responded to the critics with half-truths and evasions.

One response to the negative connotations that the Palladium and TCPA names had rapidly acquired was simply to stop using them. Palladium became the Next Generation Secure Computing Base (NGSCB) on January 25, 2003, and the TCPA became the Trusted Computing Group (TCG) on April 8, 2003.

Microsoft has continued to claim neutrality with respect to the criticisms of Palladium/NCSCB for lock-in, DRM and censorship. For the most part this has been accomplished by leaving any consideration of economic reality—both Microsoft's overwhelmingly dominant position in the applications market and the powerlessness of individual users against large organizations and network effects—out of the discussion. A few excerpts from a frequently-asked-questions list about NGSCB on Microsoft's Web site (Microsoft, 2003) will suffice to illustrate this technique.

> Q: I have heard that NGSCB will force people to run only Microsoft-approved software.

> A: This is simply not true. The nexus-aware security chip (the SSC) and other NGSCB features are not involved in the boot process of the operating system or in its decision to load an application that does not use the nexus. Because the nexus is not involved in the boot process, it cannot block an operating system or drivers or any nexus-unaware PC application from running. Only the user decides what nexus-aware applications get to run.

This is plausible only if one ignores Microsoft's overwhelming market share in office applications. As a practical matter, using a computer in today's environment includes transmitting and receiving document files created by Microsoft Word and Microsoft Excel. If future versions of those applications use NGSCP features to make their file formats inaccessible to competing applications, the competing applications will be rendered useless and the question of whether NCSCB will allow the user to run them will be moot.

> Q: What is the difference between NGSCB and DRM?
>
> A: [...] NGSCB is not DRM. The NGSCB architecture encompasses significant enhancements to the overall PC ecosystem, adding a layer of security that does not exist today. [...] A DRM system can take advantage of this environment to help ensure that content is obtained and used only in accordance with a mutually understood set of rules.

This statement is analogous to saying that a rifle creates a secure tubular structure with an opening at one end and a firing pin at the other, and that ammunition makers can take advantage of this environment to create products that accelerate projectiles to high

speeds. It fails to acknowledge that, as Microsoft's Digital Rights Management Operating System patent (United States Patent Office, 2001) makes plain, DRM is the application for which NGSCB was intended from its inception.

The same answer continues with an attempt to persuade the reader that individual users will gain as much benefit from NGSCB as content providers:

> The powerful security primitives of NGSCB offer benefits for DRM providers but, as important, they provide benefits for individual users and for service providers. NGSCB technology can ensure that a virus or other malevolent software (even embedded in the operating system) cannot observe or record the encrypted content, whether the content contains a user's personal data, a company's business records or other forms of digital content.

Along with failing to point out that the protection of the user's own data on the local system does not require the remote attestation capability upon which DRM depends, this statement ignores the fact that end users are more likely to regard "trusted" applications, such as a product registration Web page or an email client under the control of a macro virus, as a greater threat to their privacy than the increasingly rare type of virus it refers to.

Although Microsoft typically assigns the active role to conveniently nonexistent applications when it comes to unpopular potential uses of NGSCB such as lock-in, DRM and censorship, the tide flows in the opposite direction when claims are made for features that end users can be expected to find attractive. For example, a product overview document for NGSCB states: "With NGSCB, consumers will have greater control over

how their credit card numbers are used for online purchases. People will be able to store the numbers so that only an authorized program can retrieve them, and the numbers can reside on the consumer's home computer rather than in a retailer's database" (Microsoft, 2003). There is no language here about simply creating an environment where application developers might take advantage of features to do something. Instead, the claim about the handling credit card numbers, despite being as dependent upon the decisions of future application designers as any of Lucky Green's nightmare scenarios, is presented with as much certainty as a prediction that the sun will rise tomorrow.

The TCG, perhaps because Microsoft has borne the brunt of public outrage over trusted computing, has not felt as compelled as that company to defend itself against criticism. Their Web site (www.trustedcomputinggroup.org) contains only two pages that acknowledge any controversy at all over trusted computing. The first of these is the polite response, noted in the section on privacy in Chapter 2 above, to the concerns over privacy expressed in the report of the European Union Working Group (EU, 2004; TCG, 20040. The second is a link to a press article with the combative title "Trusted Computing: Maligned by Misrepresentations and Creative Fabrications" (Enderle, 2004). A reader whose interest is piqued by that title will find the article itself disappointingly bland, because it points no finger at any maligner, misrepresenter or fabricator. Instead, it identifies—and denies—two "misconceptions by the public"—that the TCG is controlled by Microsoft or the United States government, and that digital rights management is at the heart of its agenda. The first of these is best interpreted as an attempt to distance the TCG from the main storm over NGSCB and to allay suspicion in

some countries that the TCPA had been allied with Senator Hollings. As for DRM, the author deftly but unconvincingly inverts the adversarial relationship between the user and the content provider, turning the former into the demander of assurance: "Critical to this initiative is the creation of a secure repository where you can place this media and a solid trust relationship between the personal computer and the media supplier so the user can be sure the supplier is who he claims to be."

Both Microsoft and the TCG have responded to critics' charges that trusted computing would lead to lock-in, digital rights management and censorship by repeatedly protesting their innocence. It is difficult not to conclude—given the economic incentives involved, the capabilities that trusted computing would give to providers and large organizations, and the difficulty of identifying any clear and direct benefits provided to end users by those capabilities—that these companies have decided to treat the critics as a public-relations problem rather than to engage them in actual debate. This is not surprising, but it does not mean that there is no debate to be had on the subject. Trusted computing has at least one proponent who clearly sees no need to clear his arguments with any corporate communications department.

At the height of the initial furor over Palladium the British essayist Bill Thompson posted an online article (Thompson, 2002) in which he takes the position that trusted computing would improve the Internet by making it less free. Specifically, he argues that by indelibly linking every document and message transmitted over the Internet to the specific computer where it originated, trusted computing would put an end to the jurisdictional confusion that, in his view, has prevented governments from

regulating Internet use within their national territories. Without effective regulation, the Internet has been shaped by American values and turned into an instrument of American cultural imperialism, and trusted computing offers the world, and Europe especially, an opportunity to reclaim the Web from the Americans.

Thompson's strident anti-Americanism, his simplistic assessment of trusted computing's capabilities, and the equally simplistic solution that he proposes are of less interest in the present context than his willingness to oppose its critics on ideological grounds—to say not that they are mistaken but that they are wrong:

> In the mapped network we will not have the absolute freedom of speech which cyberlibertarians claim they want, but neither will we get absolute oppression, absolute free market capitalism or even absolute communism. We will instead get compromise, and regional or national variation, just as in the real world.
>
> Many will see this as a loss of freedom, but the freedom they value so much is also the freedom to act irresponsibly, to undermine civil authorities and to escape liability. It is the freedom to release viruses, abuse personal data, send unlimited spam and undermine the copyright bargain. It is not a freedom we need.
>
> It is easy to see why this approach will be resisted by US activists, of whatever political persuasion, who see the 'one world, one cyberspace' approach as a convenient way to establish an online constitutional hegemony. It will also be resisted by many of those who see any attempt to create trusted software running

> on secure processors as the network equivalent of the arrival of the black helicopters from the UN World Government Army.
>
> However, their position is untenable, because the vast majority of Internet users need and want a secure network where they can use email, look at Websites, shop, watch movies and chat to friends, and they are happy to accept that this is a regulated space just as most areas of life are.

One does not have to agree with Thompson to acknowledge that his article raises questions that the debate to date has not addressed. Are the critics of trusted computing defending a "cyberlibertarian" vision of absolute freedom? Is such a vision realistic, intellectually honest, and amenable to compromise? Is there something particularly American about it? And is it shared in some way by Microsoft and the TCG? The next chapter will examine some theoretical perspectives that offer tools that may help in making sense of these questions.

CHAPTER 4: THEORETICAL PERSPECTIVES

It is ironic, given the subject of the debate over trusted computing, that what separates the two sides is less disagreement than distrust. Users do not trust vendors to refrain from imposing restrictions on their freedom, and vendors do not trust users to refrain from undermining their property interests. It is difficult to dismiss the suspicion on either side as unjustified. At the same time, the two sides do not appear to be in sharp ideological disagreement. Despite the accusations of bad faith on both sides, it is difficult to imagine that either side does not adhere to the capitalist/libertarian ethos of individualism, market capitalism and technological progress. That may in fact be part of the problem.

One is tempted to ascribe the debate to a conflict of economic interests. In this view, computers and the Internet are following the same trajectory that other transforming technologies have followed in the past, in which an initial period of unregulated innovation has been inevitably followed by consolidation and rationalization. To borrow Barlow's (1994) perennial metaphor, this view would suggest that the electronic frontier, the abstract territory over which the proponents and opponents of trusted computing are fighting for control, is undergoing an process in which trusted computing is attempting to play the same role that barbed wire played in the taming of the American West, parceling it up and making it safe first for shopkeepers and ultimately for Wal-Mart. The frontier metaphor is rooted by now in the folklore of

computing, and it continues to appeal to both sentiment and the way many people—especially people who grew up on a steady diet of Western films and television programs—experience the Internet.

Of course the frontier view of cyberspace runs the risk of seeming pessimistic and nostalgic for the opponents of trusted computing, given the way things turned out on the old frontier. As we have seen, Benkler (2001) and Forno (2002) have insisted that cyberspace will remain wild and free because the nature of the technology upon which it is built favors new forms of information production that are based upon freedom and openness. This view does not so much repudiate the frontier metaphor as amend it by attributing a kind of incorruptible perfection to the new frontier that the old one was unable to sustain. It retains the profoundly deterministic principle that technological change is an exogenous factor in the evolution of cyberspace.

If this view were accurate, the contest over trusted computing could be expected to sort itself out through a process of natural selection, in which business models would survive or die out according to their compatibility or lack of compatibility with the environment. This chapter will introduce three writers—a legal scholar, a sociologist and a philosopher—who look at social change from broader and less deterministic perspectives. Each of them has written extensively about the role of information technology in contemporary society, and their work has yielded several insights that can open the door to both a clearer understanding of the trusted computing debate and the formulation of a strategy for moving that debate beyond its present impasse.

Lawrence Lessig

Although clearly sympathetic to the ideal of an Internet unimpeded by a draconian intellectual property rights regime, Lessig (1999) cautions against founding any hopes for it upon any assumption about the intrinsic nature of the Internet. A central idea of *Code and Other Laws of Cyberspace* is that because the Internet is a purely human artifact, its technical architecture can be expected to change over time in response to changes in the identity, needs and power of the social entities that control that architecture. Lessig points out that the Internet was originally created for the purpose of scientific collaboration, for which its open protocols and lack of security mechanisms were well suited. For the Internet to serve as a medium for commerce in the twenty-first century, he reluctantly admits, that architecture is not appropriate and will necessarily be replaced by "a far more general architecture of trust—an architecture that makes possible secure and private transactions" (p.40). Lessig argues that the shape of the new architecture—its *code*—is effectively a form of law in its power to constrain and regulate the actions of computer users, and his concern is that the making of this law is being left in the hands of private sector entities whose interests do not necessarily coincide with the public interest.

The theoretical framework that Lessig constructs for examining the ways in which user behavior is regulated in cyberspace is based on observation of analogous processes in ordinary, off-line life (p.87). Lessig identifies four constraints upon individual behavior: laws (in the traditional, narrow sense of the word), norms, market forces and the architecture of the built environment. These four regulators work simultaneously, sometimes reinforcing one another and sometimes counteracting one another. Typically

they are not equally powerful, and the distribution of power among them at a given time and place depends on a wide variety of social, political, economic and historical circumstances. In the case of the Internet at the beginning of the new century, norms are losing power, markets and laws are gaining, and architecture—code, always the most powerful regulator in the abstract and artificial environment of cyberspace—is changing from a facilitator of individual freedom into a constraint upon it. Seen from this perspective, neither the assertions by critics of trusted computing that it runs counter to the nature of the technical character of the Internet (Benkler, 2001; Forno, 2002) nor the protestations of neutrality on the part of its proponents are accurate. What is unclear in both cases is whether and to what degree the inaccuracy is rooted in naïveté or in disingenuousness.

In a later publication, Lessig (2001) revisits the idea of the interplay among the regulators of behavior in the specific context of trust. Here the crucial interplay is between norms and architecture, and the form of that interplay is the substitution of trust in technology for trust in people. Using a hypothetical online discussion group as an example, Lessig poses the question of how one member might decide whether or not to believe another member's claim to be a medical doctor. There are, he argues, two possibilities. First, one could observe the supposed doctor's behavior in the group, including interactions both with oneself and with other members, long enough and carefully enough to arrive at an informed judgment of whether that person's claim is credible. This is the process of developing trust. The alternative is to demand a credential that proves the other person is a doctor; in the online environment in question,

such a credential could be a digital certificate signed by the American Medical Association and verifiable through a public key infrastructure. Given the difficulty of evaluating the representations of himself or herself that a stranger may make online, Lessig acknowledges that in most instances the second alternative is the more rational choice: "We can thus use code, or a technical architecture, to make it so that we do not need to trust. Or more precisely, we can trust the technology rather than develop the knowledge we need to trust humans" (p. 331). This recalls Ken Thompson's (1984) article, which concluded that trusting any piece of computer software amounts to trusting all the people who participated in its creation. There is no escaping the necessity of deciding whom to trust, and the rationality of Lessig's second alternative ultimately depends upon a calculation that a system that is vouched for by large institutions that have reputations to protect and have been trusted in the past without incident by many people has a high probability of being trustworthy. This calculation is what the entire edifice of trusted computing is built upon.

Lessig argues that when reliance upon technology is substituted for reliance upon norms in a substantial number of people's everyday social transactions, a loss of social capital, in the form on an atrophying of people's ability to create and enforce norms, may be an unintended consequence: "For any particular trade-off, it might be individually more rational to substitute technology for trust, but collectively, the cost may well outweigh the benefit" (p. 331). That what is happening in this regard to the Internet is an instance of a more general phenomenon that has been at work for a very long time he readily acknowledges, but he claims that present circumstances—presumably the passage

of newly restrictive copyright legislation in the late 1990s and the initial work of the TCPA in early 2001, along with the ongoing commercialization of the Internet during those years—give it a special urgency:

> This is important now not because this is the first time we have seen an interaction between norms and technology. Obviously, technology is not new, and we have been struggling with the effects for as long as we have been struggling. But what is new is a difference of degree that matures, in my view, into a difference in kind. So plastic and so controllable is the environment of cyberspace, and so complete and pervasive will that environment become in our life, that we must with new energy focus a series of questions about how one may affect the other (p. 332).

The remainder of this chapter will introduce two writers who have looked into these questions in the broader context of the interplay of society and technology. It will identify portions of their thought that bear on the trusted computing debate. Those ideas will serve as a conceptual foundation for the discussion in Chapter 5.

David Lyon

Lyon is a sociologist who has written extensively on the phenomenon of surveillance in modern society. In *Surveillance society* (Lyon, 2001), he documents and analyzes the widespread adoption of surveillance, until recently primarily a tool used by law enforcement, other government agencies and employers to monitor criminal suspects,

recipients of benefits and employees, by commercial businesses which place entire populations under surveillance for the purpose of managing the consumption behavior of their customers. He notes that the institutional boundaries between different types of surveillance are highly porous, because of the combination of a shared emphasis on predicting people's behavior and a dependence upon a shared technological infrastructure of interconnected databases.

Contemporary surveillance places a heavy emphasis on prediction, Lyon argues, because it is first and foremost a tool for risk management. Foremost among the risks that modern businesses have had to face is uncertainty about the behavior of customers, where a miscalculation or an unlucky guess can lead to disastrous investment decisions. In the early days of market research, the collection and processing of fine-grained detail about consumer behavior was neither technically nor economically feasible, limiting businesses to coarse and often inaccurate measurements and assessments of aggregations of customers. By the 1990s however, an infrastructure was in place, spanning corporate systems, retail points of sale and commercial Web sites, that permitted the collection of detailed information about individual consumers' shopping behavior. The collection of that information is one form of surveillance, which Lyon defines as "the means whereby knowledge is produced for administering populations in relation to risk" (Lyon, 2001, p. 6).

Two of Lyon's insights into contemporary surveillance are pertinent to the subject of trusted computing. The first of these is the connection of surveillance to risk management relative to consumer behavior. It opens a perspective from which trusted

computing is revealed to be one more element, albeit a very powerful one, to be added to a pervasive pattern of surveillance that is already deeply embedded in everyday contemporary life. Viewed from that perspective, equipping a consumer's media player with digital rights management features that, as part of the process of enforcing the terms of license agreements, will report to the studio the details of that person's movie-viewing behavior, is a way to mitigate risks that go beyond piracy. Combined with similar information, collected from a wide variety of sources, about that same consumer's tastes and purchases in other product categories, the information collected by the media viewer could be used to reduce the risks in countless corporate investment decisions, from casting a certain star in a new film to adding a item to the menu of a local restaurant. At the same time, the same information can be used to manage, and thereby reduce the risk inherent in, that same individual's future behavior, by creating and presenting to that person advertising specifically tailored not only to her tastes but also to her present situation. Taken to its limit—a limit that trusted computing seems to place tantalizingly close to being within reach—this type of consumer surveillance would constitute a type of perfection, where risk would vanish altogether and unsold inventories would be a thing of the past.

The second of Lyon's insights is that much of the surveillance that is done in contemporary society is done with the approval and willing participation of the public.

> Surveillance always has two faces, and part of the problem of convincing people about the more worrisome and unsocial aspects is that they appear merely as the price one pays for the speed, safety and security apparently offered by the other

> 'face'. Needless to say, those government departments and corporations that stand to gain from surveillance are in a good position to make their case (p. 136).

Lyon ascribes the public's acquiescence in surveillance in part to "the hegemonic power suffusing many contemporary technologically advanced societies" (p. 136), which he describes, like surveillance itself, as a coin with two faces. The first face is the simple willingness to accept surveillance as the price of the conveniences and security of modern society. The second and less obvious face is "a widespread assumption that rights to privacy comprise the appropriate language for questioning surveillance when necessary" (p. 136). When one objects to surveillance on the grounds that it violates privacy, one has already, in Lyon's opinion, agreed to look at the world in the same way that led to surveillance in the first place:

> Privacy talk [...] is still part of the hegemonic system of consent to the dominant liberal culture of law and the establishment. It will not go beyond these to question the very worldviews and power bases of those who have access to the surveillance switches (p. 137).

Lyon's explanation of this hegemonic system is cursory at best. He is making an important point here, which will be revisited in the next section. For the time being, however, it will suffice to note that the controversy over trusted computing follows the pattern Lyon describes, at least with respect to the behavior of Microsoft and the TCG. Those organizations are making the case for trusted computing on the basis of apparent security and safety. As noted in Chapter 2, the only criticism to date that either organization has taken seriously enough to modify its specification (EU, 2004) was

phrased in terms of privacy and came from a quasi-government body. On the other side of the debate there appears to be a wider range of opinion, ranging from polite calls for accommodating privacy concerns (EU, 2004; Arbaugh, 2002) to opposition to the entire project (Anderson, 2003; Forno, 2002; Green, 2002b). To evaluate the extent to which these critics share an ideology with each other and even with Microsoft and the TCG, a closer and more rigorous examination of the hegemonic system mentioned by Lyon will be necessary. That is the goal of the next section.

Albert Borgmann

Placing our trust in technology extends beyond trusting that a particular mechanism will do what we require it to do. On a more general level, it means looking at the world in a way that presents our wants, needs and cares as problems that are amenable to technological solutions.

The philosopher Albert Borgmann created in *Technology and the Character of Contemporary Life* (1984) an analytical framework for thinking about how technology and ideology interact. At the heart of Borgmann's analysis of technology is the distinction between technological *devices* and pretechnological *things*. Both of these terms have very specific meanings in Borgmann's framework. He explains the distinction by comparing a wood-burning stove to a central heating system. The stove is an example of a thing:

> A thing, in the sense in which I want to use the word here, is inseparable from its context, namely, its world, and from our commerce with the thing and its world, namely, engagement. The experience of a thing is always and also a bodily and social engagement with the thing's world. In calling forth a manifold engagement, a thing necessarily provides more than one commodity. Thus a stove used to furnish more than mere warmth. It was a focus, a hearth, a place that gathered the work and leisure of a family and gave the house a center (p. 41).

The central heating system is a technological device:

> A device such as a central heating plant procures mere warmth and disburdens us of all other elements. These are taken over by the machinery of the device. The machinery makes no demands on our skill, strength, or attention, and it is less demanding the less it makes its presence felt. In the progress of technology, the machinery of a device has therefore a tendency to become concealed or to shrink. Of all the physical properties of a device, those alone are crucial and prominent which constitute the commodity that the device procures. Informally speaking, the commodity of a device is "what a device is there for" (p. 42).

The point that Borgmann is making in this distinction is that there is a loss as well as a gain when a thing is replaced by a device. In some cases what is lost is simply drudgery, but in others the burden that the device removes includes things of value, such as the creation of meaning, the development or exercise of skill, and participation in a rich social context. Similarly, the value of the commodity procured by a device can vary, ranging from the cure of a deadly disease to a moment of frivolous pleasure. Given the

variability in the terms of the trade-off, one might expect that people would, when faced with the choice between a thing and a device, weigh the alternatives carefully and choose the device only when the commodity it supplies is valuable enough to compensate for what is lost in giving up the thing. Such careful and deliberate choices, Borgmann is quick to point out, are more the exception than the rule in modern society.

Daily life in the modern world presents people with an endless succession of choices between devices and things: a wood burning fireplace or a gas log, a frozen dinner or a home-cooked meal, a stereo set or a violin. It is Borgmann's contention that the way modern societies are organized favors the choice of devices over things, often so overwhelmingly that engagement with things becomes a practical impossibility. This predisposition toward technology operates in our collective decisions as well as in our individual choices, and it has been played a major role in the shaping of modern institutions in government, education and commerce. Borgmann has named it the *device paradigm*, and he stresses its importance and centrality to his work in the opening paragraph of his first chapter:

> I propose to show that there is a characteristic and constraining pattern to the entire fabric of our lives. This pattern is visible first and most of all in the countless and inconspicuous objects and procedures of daily life in a technological society. It is concrete in its manifestations, closest to our existence, and pervasive in its extent. The rise and rule of this pattern I consider to be the most consequential event of the modern period (p. 3).

Borgmann's analysis of the device paradigm reveals it to be embedded in the other institutions of modernity, including capitalism, individualism and liberal democracy. This is the hegemonic system that Lyon (2001) referred to in the final citation of the previous section.

Borgmann traces the device paradigm to the beginning of the modern period, when the promise of technology was first formulated by Bacon and Descartes. From that time forward, the proponents of new technology in each generation have reiterated the promise that technology was on the verge of ushering in a new era of freedom, prosperity and happiness. In the early years of modernity, when the strides taken by technology toward liberating mankind from the threats of starvation and disease were dramatic and unprecedented, the promise of technology was easy to believe. By the time that humanity had gained enough experience with technology to have become skeptical about the final fulfillment of its promise, Borgmann argues, the device paradigm had acquired the character of an ideology. Institutions, value systems and habits of thought had grown up around it, to the point where the natural, common-sense answer to the shortcomings of any present technology was the improved and perfection technology just over the horizon.

The pervasive, taken-for-granted character of the device paradigm extends beyond daily life and into theoretical discourse as well. The result, Borgmann argues, is that most contemporary debates about technology are constrained by the assumptions built into the device paradigm. This manifests itself in two ways that are pertinent to the debate over trusted computing. First, the hopes for new technologies are typically

exaggerated to utopian proportions as each new "breakthrough" is looked to as the one to finally usher in the new era of ease and prosperity. Second, values and concerns to which the device paradigm is indifferent or blind, such as Lessig's (2001) concern about social capital, are ignored, with the result that every question is formulated in a way that is conducive to a technological answer. These two characteristics of the device paradigm can be used to explain the failure so far of the debate over trusted computing to address substantive issues.[1]

First, there is the issue of exaggerated hopes. Throughout its history the computer has been a magnet for visionaries. Even before the invention of the transistor, scientists such as Bush (1945) were imagining a future in which machines would make vast stores of information available to their users. In the early mainframe era, when only large organizations could afford computers, the visions that accompanied and drove information technology were largely confined to institutional settings, in which the new machines were expected to deliver organizations from the inefficiency and risk inherent in dependence upon slow and error-prone clerical workers. Later, as terminals and eventually complete computers became cheap enough for individuals to afford, a second set of visions came along. Inspired by the countercultural rebellion of the 1960s, these dreams were more personal and antiauthoritarian than their predecessors. Instead of organizations freed from inefficiency or office workers freed from tedium, the visions

[1] It is possible that the silence on the part of the part of the proponents in industry and government of trusted computing is strategic, like the refusal of a popular incumbent political candidate to debate an underdog challenger. However, this in no way precludes the very likely possibility that the people behind NGSCB and TCG sincerely believe that trusted computing will make the world a better place, and that their public statements reflect their actual thinking on the subject.

associated with personal computers featured individuals freed from social constraints of all types.

Both of these visions, of perfect control and of perfect freedom, are still alive and influential today. Although they may seem mutually antithetical, and are locked in battle in the present debate over trusted computing, it is a fact that they have managed to coexist for over two decades without getting in each other's way. The software industry itself, which today is animated by the vision of control, has on the rare occasions where the visions have conflicted, most notably the controversies over cryptography and access to pornography by children, generally aligned itself with the libertarian vision. This compatibility over time suggests that the two visions may have more in common than their present conflict might lead a casual observer to suspect. It may even suggest that the two visions are really one and that the present conflict is a manifestation of an underlying contradiction within that single vision. This hypothesis in consistent with Borgmann's theory of the device paradigm and bears further examination.

For any individual agent, perfect freedom and perfect control are inseparable. If that agent's actions are to be unconstrained by the actions of others, he or she must be able to control the actions of others. Control, at least when considered as an absolute, is simply another way of looking at freedom. Whether conceived in terms of freedom or of control, the vision of technological perfection becomes self-contradictory as soon as the possibility of multiple agents is admitted, because perfect freedom can never be the possession of more than one agent. For the person who is clinging to such a vision, this presents a problem that is most easily solved by refusing to take seriously the possibility

of conflict. This refusal takes the form of ignoring or delegitimizing opposing points of view. In the utopia of perfect control, only a thief would object to digital rights management, and in the utopia of perfect freedom, only a greedy and desperate corporation would attempt to interfere with an individual's control of his or her own computer. Thus on both sides of the controversy over trusted computing, utopian visions inspired by the device paradigm militate against both honest debate and compromise.

The second characteristic of the device paradigm that is at work in the trusted computing debate is it tendency to blind both parties to the non-technological dimensions of the subject at hand. The public statements of Microsoft and the TCG consistently present their proposals as straightforward, stand-alone technical solutions to unambiguous problems. Users want security, and trusted computers will build it into the system for them. Similarly, the opposition believes that maintaining the Internet's open protocols and making software non-proprietary will lead naturally to the evolution of new and better ways to use the Internet. Both of these view are grounded in technological determinism, both in their disregard for the role of non-technological forces in the shaping of society and in their tacit appeal to a technological destiny that is waiting to come into being.

CHAPTER 5: DISCUSSION

The debate over trusted computing, no matter how it is conducted, will necessarily end by deciding where the limits of various freedoms will be drawn by some form of regulation. These include the freedom of users to use downloaded digital materials as they see fit, the freedom of software vendors to control access to their file formats, the freedom of content providers to license their products instead of selling them, and the freedom of computer owners to choose the software they wish to use, and other freedoms too numerous to list. The only prediction that can be made with certainty about where those limits will be drawn is that the result will necessarily fall short of anyone's idea of perfection.

It is, however, possible to identify the most probable outcome. Five years ago, Lessig predicted that if then-present trends continued, the architecture of the Internet would undergo a radical transformation:

> [...] the invisible hand of cyberspace is building an architecture that is quite the opposite of what it was at cyberspace's birth. The invisible hand, through commerce, is constructing an architecture that perfects control—an architecture that makes possible highly effective regulation (1999, p. 6).

His use of the word *perfects* was something of an overstatement ("seeks to perfect" would have been more accurate), but overall his prophecy has proven accurate. Trusted computing has been proposed by powerful commercial interests, opposition to it has been

unfocused and ineffectual, and plans are in place for a new generation of computers and operating systems incorporating NGCSB to hit the market in the near future.

As noted in the previous chapter, there are structural constraints built into technological thinking that make it difficult to identify and debate many of the issues implicit in the prospect of trusted computing. The present chapter will present the outlines of the debate that could, and should, take place if the fog of self-deception and myopia that surrounded the present debate were lifted. That debate would focus on the two primary problems that trusted computing purports to address, digital piracy and the vulnerability of systems to software attacks. For both of these subjects, the principles of the debate are the same: to acknowledge the legitimacy of the concerns of both sides, to avoid deterministic thinking, and to consider the problem at hand in its full range of dimensions.

Digital Piracy

Digital piracy is a serious problem. Producing a feature film, for example, typically requires the investment of many millions of dollars, and the value of that investment can be seriously damaged if large numbers of people view pirated copies of that film instead of paying to see it legitimately. Conversely, heavy-handed and draconian measures that place the entire population under surveillance and suspicion for the sake of preventing piracy impose costs on society that need to be recognized and taken seriously. That being said, there are two questions that need to be debated with

regard to the distribution of digital content: how intensively should that distribution be regulated, and what are the proper regulators to use?

How intensively the distribution of digital content should be regulated is essentially an economic question. Answering it will require a realistic assessment of the risk that distributing digital media content over the Internet imposes on content providers under the present regulatory scheme, along with an assessment of the costs, both direct and indirect, of a number of alternative regulatory schemes. The indirect costs to be considered cover a wide range of social costs, including inconvenience to individual consumers, the legal expenses of complex licensing schemes, slowing of innovation, and the effect on public respect for law of treating consumers as potential criminals. The purpose of these assessments is to fit the solution to the problem and to ensure that the costs imposed by the solution are paid by those who reap its benefits.

Identifying the proper regulators for digital distribution is a task where assigning each present or possible regulator to one of Lessig's (1999) classes—laws, norms, markets and code—is a useful place to start. These categories are useful both because they stimulate the imagination to recognize a wide range of possibilities and because the regulators within each category tend to have characteristics in common in terms of their social impact.

In theory at least, laws are the form of regulation most amenable to democratic oversight. Even when bad laws are passed by ignorant or corrupt legislators, it is less costly to amend such laws than, for example, to rework the architecture of a large infrastructure. Until the passage of the Digital Millennium Copyright Act (1998), laws

were the dominant regulator of the distribution of information. Aspects of laws that are relevant to distribution of digital information include the difficulty of identifying which authority has jurisdiction in a specific case, the lack of uniformity among different jurisdictions, the presumption of innocence, the emphasis on sanctioning violations rather than preventing them, the cost of legal administration, the danger of classifying large numbers of otherwise law-abiding citizens as criminals, and the comparative merits of copyright protection and licensing.

Norms are the oldest most powerful form of social regulation. They are difficult and expensive to create, and they are capable of taking a shape contrary to the wishes of policy makers, but once established they perpetuate and enforce themselves with a minimum of direct costs. At the present time the music industry and educational institutions throughout the United States are engaged in a difficult effort to change the existing norms among young people who do not regard the sharing of MP3 files as improper behavior. As Lessig (2001) has pointed out, norms can be a valuable cultural asset, and when they are allowed to atrophy through the substitution of other regulators a significant social loss can result. Norms are most powerful for well-defined and homogeneous communities, and their power can rapidly diminish when their domain of application experiences exponential growth. This is illustrated by the fate of “netiquette”, which in the earliest days of the Internet and online communities such as the WELL (Rheingold, 1993), was a much stronger regulator of online behavior than it is today.

Markets regulate behavior by appealing to the self-interest of buyers and sellers. Ideally, market transactions are freely entered into by both parties, making markets less

coercive in theory than other regulators. It is for this reason that some scholars (Benkler, 2001; Davis, 2001) have recommended that content providers develop business models that would discourage digital piracy by making the value proposition of legitimately purchasing content more attractive to the consumer than downloading it for free from the Internet. Typically, these models involve distributing the content in a basic form for free while offering supplemental services for sale such as documentation or upgrades in audio and/or video quality. Whether such models could succeed would depend of course on how susceptible the supplements would themselves be to piracy.

Code is the regulator of choice for the proponents of trusted computing. It offers advantages of simplicity. Changes in architecture are relatively straightforward to make in comparison with other regulators, especially in highly plastic products like software and integrated circuits. This is true both technically and in business terms when the products in question are controlled by a small number of allied firms. At the same time, once regulation is embedded in code, compliance is guaranteed to the extent that the code changes have succeeded in making noncompliance impossible. The disadvantages of code as a regulator also flow from its simplicity. Because it is simple, the power to regulate through code—and the power to decide how, where and whom to regulate—is highly concentrated in organizations that answer to their shareholders instead of to the public. Furthermore, those organizations are populated by technologists who, as the responses to critics by Microsoft and the TCG have demonstrated, are at best dimly aware of their role as makers of social policy. Because regulation by code requires no enforcement mechanisms beyond the code itself, anyone who is subject to such

regulation is singularly powerless, with no alternatives or avenues of appeal, redress or negotiation. Finally, regulation through code is a dangerously blunt instrument. Rules embedded in code tend to be of the one-size-fits-all variety, designed to handle what their creators judge to be the most typical situations and insensitive to contextual information that other regulators would take into account as a matter of course. A clear example of this is the inability of digital rights management software to determine whether a user's attempt to copy a section of a protected document is for a purpose protected by the fair use doctrine. This insensitivity opens regulation by code up to an unlimited range of potential abuses and unintended consequences. For all these reasons, regulation through code is very dangerous and should be used both with extreme caution and only as a last resort.

Although this discussion has separated for clarity of exposition the questions of intensiveness and choice of means, these two questions are clearly interdependent. The first question cannot be answered without understanding the alternatives that are developed in answering the second, and a definitive answer to the second question depends on answering the first. In practice, a debate would answer both questions simultaneously in an iterative process that would gradually narrow the set of alternatives under consideration.

Vulnerability to Software Attacks

Viruses, worms, Trojan horses and other forms of malicious software differ from digital piracy in two crucial respects that have shaped the computing community's response to them. The harm they cause is not restricted to a single industry, and they tend to manifest themselves as discrete and disruptive events rather than as a slow leakage of assets that can be charged to the cost of doing business. Because of this, many people in industry, in universities and in government have been working for many years on developing solutions to the problem of software attacks. Tools and techniques have been developed, including firewalls, intrusion detection systems, virus checkers and virtual private networks. An infrastructure has been put in place for gathering and distributing information about vulnerabilities and new attacks. In this context, alternatives to trusted computing are, to a large but not universal extent, already present and in daily use. Many people in field believe that the challenge facing information security today is more managerial than technical, that adequate tools exist and what is needed is the wisdom and will to use them effectively. Given all this, it is not clear what new and needed capability trusted computing adds to the security toolkit.

On the down side, the same dangers that were found in the previous section to be associated with the use of code to regulate distribution of content are present with respect to the use of trusted computing as a security mechanism. In fact they are dangers in large part precisely because it is impossible to say for what purpose trusted computing would ultimately be or not be used. In the context of security, the undesirable consequences of depending upon trusted computing could include the unnecessary restriction of legitimate

user of computers by their owners, the consolidation of market power in the hands of a small number of firms, the use of trusted computing features by criminals of all varieties to evade prosecution, the diversion of programming resources in software firms away from the fixing of architectural flaws that make software vulnerable to attack in the first place, and the creation of a false sense of security vis-à-vis threats, such as macro viruses, against which trusted computing offers no protection.

The foregoing is not meant to suggest that present approaches to software security are beyond improvement. Market forces and laws in particular have not been used effectively to counter software attacks. Denning's (1993) proposal to allow the market to define what security features a system should have was doomed by its dependence upon the assumption that the software market was efficient, and more recent calls from industry and government to let the market decide about security, such as *The National Strategy to Secure Cyberspace* (2003), are not really about markets at all, but rather rhetorical attempts to brand any proposal to legally regulate the behavior of software companies as an affront to free enterprise. In fact, legal changes that would, for example, expose software firms to liability for defects in their products could be more accurately described as pro-market, because their effect would be to make the market more efficient by internalizing the costs of those defects.

CONCLUSION

This thesis has described the historical and present context in which the proposals of ubiquitous trusted computing by Microsoft and the Trusted Computer Group have been made. It has documented the controversy that has ensued, and offered an explanation for the absence so far of a productive and substantive debate on the merits of the proposed systems. Finally, it has attempted to identify the issues, questions and concerns that a proper debate should raise for discussion. In conclusion, and in the hope that it is still not too late for trusted computing to be debated and not imposed by fiat, it offers three general recommendations to all participants in the debate.

First, be honest. Acknowledge that computers and the Internet serve a wide variety of legitimate interests, and that those interests can and will conflict at times. When speaking of people's freedom of choice in the market, recognize that choice is constrained by network effects and the power of incumbents.

Second, be realistic. The proposals under consideration need to be judged for their usefulness in solving problems today and in the near future, and not for a contribution they might make to fulfilling a destiny.

Finally, be open-minded. The debate about trusted computing is not a purely technical discussion. It is in fact a legislative debate about how society should regulate the use of computer technology. As such it must consider all the tools that society has available to it, as well as all the interests that society takes in the outcome.

REFERENCES

"AARG!Anonymous" (2002, August 9). Re: Thanks, Lucky, for helping to kill gnutella. Message posted to Cryptography Mailing List, archived at http://www.mail-archive.com/cryptography@wasabisystems.com/msg02533.html

Anderson, R. (2002a, June 28). "Palladium" and TCPA. Message posted to UK Cryptography Policy Discussion Group, archived at http://www.chiark.greenend.org.uk/pipermail/ukcrypto/2002-June/019444.html

Anderson, R. (2002b, June 29). "Palladium" and TCPA. Message posted to UK Cryptography Policy Discussion Group, archived at http://www.chiark.greenend.org.uk/pipermail/ukcrypto/2002-June/019463.html

Anderson, R. (2002c, July). Trusted Computing Frequently Asked Questions, Version 1.0. Retrieved April 6, 2004, from http://www.cl.cam.ac.uk/~rja14/tcpa-faq-1.0.html

Anderson, R. (2003, August), 'Trusted Computing' Frequently Asked Questions, Version 1.1. Retrieved April 6, 2004, from http://www.cl.cam.ac.uk/~rja14/tcpa-faq.html

Arbaugh, W. A. (2002, August). Improving the TCPA Specification. *IEEE Computer, 35* (8), 77-79.

Barlow, J. P. (1994). Jack in, Young Pioneer. Retrieved April 6, 2004, from http://www.eff.org/Publications/John_Perry_Barlow/HTML/jack_in_young_pioneer.html

Barlow, J. P. (1996). A Declaration of the Independence of Cyberspace. Retrieved April 6, 2004, from http://www.eff.org/~barlow/Declaration-Final.html

"Bear" (2002, August 14). Re: Overcoming the potential downside of TCPA. Message posted to Cryptography Mailing List, archived at http://www.mail-archive.com/cryptography@wasabisystems.com/msg02611.html

Benkler, Y. (2001, February). The Battle Over the Institutional Ecosystem in the Digital Environment. *Communications of the ACM, 44*(2), 84-90.

Borgmann, A. (1984). *Technology and the Character of Contemporary Life*. Chicago: University of Chicago Press.

Boyle, J. (2002, Spring), Fencing Off Ideas: Enclosure and the Disappearance of the Public Domain. *Daedalus*, 13-25.

Bush, V. (1945, July). As We May Think. *The Atlantic Monthly, 176*(1), 101-108.

CBDTPA (2002). S. 2048, Consumer Broadband and Digital Television Promotion Act. Introduced March 21, 2002. Retrieved April 6, 2004, from http://www.techlawjournal.com/cong107/copyright/hollings/s2048is.asp

NIST (2004). ICAT Metabase. Washington, DC: National Institute of Standards and Technology. Retrieved April 27, 2004, from http://icat.nist.gov/icat.cfm

Disney Lobbying for Copyright Extension No Mickey Mouse Effort. (1998, October 17). *Chicago Tribune*. Retrieved April 6, 2002, from http://homepages.law.asu.edu/~dkarjala/OpposingCopyrightExtension/commentary/ChiTrib10-17-98.html

Coursey, D. (2002, July 2). Why we can't trust Microsoft's "trustworthy" OS. *ZDNet.com*. Retrieved April 6, 2004, from http://zdnet.com.com/2100-1107-941111.html

Davis, R. (2001, February). The Digital Dilemma. *Communications of the ACM, 44*(2), 77-83.

Denning, D. (1993). A New Paradigm for Trusted Systems. In xxxx (Ed.), *Proceedings on the 1992-1993 workshop on New security paradigms* (pp.36-41). Little Compton, RI: ACM Press.

Department of Defense (1985). *Trusted Computer System Evaluation Criteria*. DoD Doc. 5200.28-STD.

Digital Millennium Copyright Act (DMCA) of 1998, Pub. L. 105-304, 112 Stat. 2860 (1998).

Eldred v. Ashcroft (2003). United States Supreme Court, 01-618. Retrieved April 6, 2004, from http://www.supremecourtus.gov/opinions/02pdf/01-618.pdf

Enderle, R. (2004, February 5). Trusted Computing: Maligned by Misrepresentations and Creative Fabrications. *Storage Pipeline*. Retrieved April 6, 2004, from http://subscriber.acumeninfo.com/uploads2/0/A/0AFC7F759F71BF2C4E80FAE3CFD19DCF/1079105225453/SOURCE/4643tsg.html

EU (2004, January 23). Article 29 Data Protection Working Party Working Document on Trusted Computing Platforms and in particular on the work done by the Trusted Computing Group (TCG group). Retrieved April 6, 2004, from http://europa.eu.int/comm/internal_market/privacy/docs/wpdocs/2004/wp86_en.pdf

Felsenstein, L. (1993, May). The Commons of Information. *Dr. Dobbs' Journal*.

Forno, R. (2002). Hollings, Valenti and the American Techniban (Operation ENDURING VALENTI). Retrieved April 6, 2004, from http://research.yale.edu/lawmeme/modules.php?name=News&file=article&sid=159

Gates, W. (2002, July 18). Executive E-mail: Trustworthy Computing. Retrieved April 6, 2004, from http://www.microsoft.com/mscorp/execmail/2002/07-18twc.asp

Green, L. (2002a, June 27). Two additional TCPA/Palladium plays. Message posted to Cryptography Mailing List, archived at http://www.mail-archive.com/cryptography@wasabisystems.com/msg02246.html

Green, L. (2002b, August). TCPA: the mother(board) of all Big Brothers. Presentation at DefCon X. Retrieved April 6, 2004, from http://www.cypherpunks.to/

Hollings, Ernest (2002), Statement by Senator Ernest F. Hollings on the Introduction of "The Consumer Broadband and Digital Television Act of 2002". Retrieved April 6, 2004 from http://www.politechbot.com/docs/cbdtpa/hollings.cbdtpa.release.032102.html

Jøsang, A. (1997). The right type of trust for distributed systems. In *Proceedings of the 1996 workshop on New security paradigms* (pp. 119-131). Lake Arrowhead, CA: ACM Press.

Krebs, B. (2003, February 14). A Short History of Computer Viruses and Attacks. *Washington Post*. Retrieved April 6, 2004, from http://www.washingtonpost.com/ac2/wp-dyn/A50636-2002Jun26?start=15&per=18

Landwehr, C. (1993). How Far Can You Trust A Computer? In Gorski, J. (Ed.), SAFECOMP'93, *Proceedings of the 12th international conference on Computer Safety, Reliability and Security* (pp. 313-325). Poznan-Kierkrz, Poland: Springer-Verlag.

Lessig, L. (1999). *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lessig, L. (2001, April). Preface to a Conference on Trust. *Boston University Law Review, 81* (2),329-332.

Lessig, L. (2002a). *The Future of Ideas* (Rev. Ed.) New York: Vintage Books.

Lessig, L. (2002b, September 10). Anti-Trusting Microsoft. *RedHerring.com*. Retrieved April 6, 2004, from http://www.lessig.org/content/columns/red1.pdf

Levy, S. (2001, July 1). The Big Secret. *Newsweek*, 48-50.

Licklider, J. C. R. (1960, March). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics, HFE-1*, 4-11.

Loney, M. (2002, June 27). Who trusts Microsoft's Palladium? Not me. *ZDNet.com*. Retrieved April 6, 2004, from http://zdnet.com.com/2102-1107-939817.html

Lyon, D. (2001). *Surveillance society*. Buckingham: Open University Press.

Manjoo, F. (2002, July 11). Can we trust Microsoft's Palladium? *Salon.com*. Retrieved April 6, 2004, from http://archive.salon.com/tech/feature/2002/07/11/palladium/

McAllister, W. (1998, October 15). A Capital Way to Stop a Headache. *Washington Post*, p. A21. Retrieved April 6, 2004, from http://www.public.asu.edu/%7Edkarjala/commentary/WashPost10-15-98.html

McCullagh, D. (2002, March 30). Hollings Howls Will Have to Wait. *Wired News*. Retrieved April 6, 2004, from http://www.wired.com/news/politics/0,1283,51425,00.html

Microsoft (2002a, August). Microsoft “Palladium”: A Business Overview. Retrieved April 6, 2004, from http://www.microsoft.com/presspass/features/2002/jul02/0724palladiumwp.asp

Microsoft (2002b, August). Microsoft “Palladium” Initiative Technical FAQ. Archived on http://www.napolifirewall.com/MicrosoftPalladium.htm

Microsoft (2003a, April). The Next-Generation Secure Computing Base: An Overview. Retrieved April 6, 2004, from http://www.microsoft.com/resources/ngscb/NGSCB_overview.mspx

Microsoft (2003b, July). Microsoft Next-Generation Secure Computing Base - Technical FAQ”, July 2003, Retrieved April 6, 2004, from http://www.microsoft.com/technet/security/news/ngscb.mspx

Morrissey, B. (2002, June 28). Is Microsoft's Palladium a Trojan Horse? *Internetnews.com*. Retrieved April 6, 2004, from http://www.internetnews.com/xSP/article.php/1378731

Nelson, T. H. (1987). *Computer lib. Dream machines* (Rev. ed.). Redmond, WA: Tempus Books of Microsoft Press

National Strategy to Secure Cyberspace (2003). Washington, DC: The White House. Retrieved April 6, 2004 from http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf

NIST (2004). ICAT Metabase. Washington, DC: National Institute of Standards and Technology. Retrieved April 27, 2004, from http://icat.nist.gov/icat.cfm

No Electronic Theft (NET) Act of 1997, Pub. L. No. 105-147, 111 Stat. 2678 (1997).

NRC (2002), National Research Council Computer, Science and Telecommunications Board, Division on Engineering and Physical Sciences, *Cybersecurity Today and Tomorrow*. Washington, DC: National Academy Press.

Olsen, F. (2003, February 21). Control Issues: Microsoft's plan to improve computer security could set off fight over use of online materials. *The Chronicle of Higher Education*. Retrieved April 6, 2004, from http://chronicle.com/free/v49/i24/24a02701.htm

Rheingold, H. (1993). *The Virtual Community: homesteading on the electronic frontier* (1st HarperPerennial ed.). New York: HarperPerennial.

Roszak, T. (1994). *The Cult of Information: A Neo-Luddite Treatise on High Tech, Artificial Intelligence, and the True Art of Thinking* (2nd ed.). University of California Press.

Samuelson, P. (2000). Towards More Sensible Anti-circumvention Regulations. *Proceedings of Financial Cryptography 2000 Conference*. Retrieved April 6, 2004, from http://www.sims.berkeley.edu/~pam/papers/fincrypt2.pdf

Schneier, B. (2002, August 15), Palladium and the TCPA. *Crypto-Gram Newsletter*. Retrieved April 6, 2004, from http://www.schneier.com/crypto-gram-0208.html

Schoen, S. (2002a, July 3). Web log. Retrieved April 6, 2004, from http://vitanuova.loyalty.org/2002-07-03.html

Schoen, S. (2002b, August 9). Web log. Retrieved April 6, 2004, from http://vitanuova.loyalty.org/2002-08-09.html

Schoen, S. (2003). Trusted Computing: Promise and Risk. *Electronic Frontier Foundation*. Retrieved April 6, 2004, from http://www.eff.org/Infra/trusted_computing/20031001_tc.php

Simons, B. and E. H. Spafford. (2002, March 29), Letter to Senator Ernest F. Hollings. Retrieved April 6, 2004, from http://www.acm.org/usacm/Legislation/Holllings_S2048.htm

Sonny Bono Copyright Term Extension Act (CTEA) of 1998, Pub. L. 105-298, 112 Stat. 2827 (1998).

Stallman, R. (2002, October 21), Can you trust your computer? *NewsForge*. Retrieved April 6, 2004, from http://www.newsforge.com/business/02/10/21/1449250.shtml?tid=19

TCG (2004, February 6). TCG talks about the recent European Union report on TCG and its specifications. Retrieved April 6, 2004, from https://www.trustedcomputinggroup.org/press/feb_6_art_29_report_QA.pdf

TCPA (2000a, January), Building A Foundation of Trust in the PC. Originally retrieved from TCPA Web site, no longer available. On file with the author.

TCPA (2000b, August), TCPA Security and Internet Business: Vital Issues for IT. Originally retrieved from TCPA Web site, no longer available. On file with the author.

Thompson, K. (1984, August). Reflections on Trusting Trust. *Communications of the ACM*, 761-763

Thompson, W. (2002, August 9). Damn the Constitution: Europe must take back the Web. *The Register*. Retrieved on April 6, 2004, from http://www.theregister.co.uk/2002/08/09/damn_the_constitution_europe_must/

Udell, J. (2002, July 9), Control Your Identity or Microsoft and Intel Will. *O'Reilly Webservice.xml.com*. Retrieved April 6, 2004, from http://webservices.xml.com/pub/a/ws/2002/07/09/udell.html

United States House of Representatives Subcommittee on Courts and Intellectual Property (1997, September 11). Hearing on Copyright Piracy, And H.R. 2265, The No Electronic Theft (NET) Act. Retrieved April 6, 2004, from http://commdocs.house.gov/committees/judiciary/hju48724.000/hju48724_0f.htm

United States Copyright Office (2004). Copyright Legislation Archive. Retrieved April 6, 2004, from http://www.copyright.gov/legislation/archive/

United States Patent Office (2001, December 11). United States Patent 6,330,670. Retrieved April 6, 2004, from http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=/netahtml/srchnum.htm&r=1&f=G&l=50&s1=6330670.WKU.&OS=PN/6330670&RS=PN/6330670

USENIX (2002). Conference Reports, 11th USENIX Security Symposium. Retrieved April 6, 2004, from http://www.usenix.org/events/sec02/confrpts.pdf

Working Group on Intellectual Property Rights. (1995, September). *Intellectual Property and the National Information Infrastructure, The Report of the Working Group on Intellectual Property Rights*. Washington, DC: U. S. Patent and Trademark Office. Retrieved April 6, 2004, from http://www.uspto.gov/web/offices/com/doc/ipnii/