Testimony of Edward W. Felten
Associate Professor of Computer Science at Princeton University and Director of the Secure Internet Programming Laboratory

Submitted to:

The Senate Judiciary Committee Hearing on
"Competition, Innovation, and Public Policy in the Digital Age: Is the Marketplace Working to Protect Digital Creative Works?"

March 11, 2002

To the Distinguished Members of the Senate Judiciary Committee:

I am writing to express my concern about proposals for laws that mandate the inclusion of copy protection technologies in computer hardware and software. Most of the discussion of these proposals has focused on their effect on fair use and the rights of consumers, issues that deserve serious consideration. However, that is not my topic here. I write instead to raise another issue that is equally important but has received much less attention: the technical effect of the proposed legislation.

I write as an expert on computer security and copy protection technology. I am an Associate Professor of Computer Science at Princeton University, and Director of Princeton's Secure Internet Programming Laboratory. (I am currently on sabbatical leave at the Center for Internet and Society at Stanford Law School.) I have published more than fifty research papers and two books, and my research has been covered widely in the national press. In addition to my service on corporate advisory boards, I serve on the Information Science and Technology (ISAT) advisory board of the Defense Advanced Research Projects Agency. I am co-chair of an ISAT study on "Reconciling Security with Privacy," and am a member of the National Research Council's study group on "Fundamentals of Computer Science." I have also served as the primary computer science expert witness for the Department of Justice in the Microsoft antitrust case. I believe that the views I express here are mainstream ones among independent computer security experts.

The content industry, including the movie studios and the record companies, is afraid of copyright infringement on the Internet, and with good reason. Although there are disagreements about exactly what constitutes infringement, all serious participants in the debate agree that large-scale infringement, especially the unauthorized redistribution of copyrighted movies or music, is a serious problem that will require a serious effort to solve.

In their zeal to respond to this threat, the content industry proposes the creation of a standard for copy protection, a standard that is intended to prevent the unauthorized copying of copyrighted works. The key word here is "intended," for clearly intentions are not enough. The standard must work – it must actually prevent would-be infringers from copying.

All indications are that it will not work. To date there is little if any scientific evidence to indicate that a technology of the sort envisioned by the content industry could actually prevent piracy. The consensus among independent experts, including me, is that strong copy protection (protection that a moderately skilled person expending moderate effort cannot break) simply is not possible on general-purpose computers such as PCs. A strong copy protection scheme for PCs is as implausible to many experts as a perpetual motion machine.

While copy protection might be workable in a world with "dumb" single-purpose media players like VCRs, it is fundamentally incompatible with "smart" general-purpose technologies such as PCs and the Internet. When I say that these technologies are "general-purpose," what I mean is that they are able to perform powerful operations on data, without needing to understand everything about that data. For example, the telephone system is a general-purpose technology, because it can carry a conversation between two faraway people, and it can do this without needing to understand what those people are talking about. The telephone is indispensable precisely because you can use it to talk about any topic whatsoever, and because it transmits faithfully every pause, inflection, and nuance in the speakers' voices; and it is feasible to build a flexible, inexpensive, and easy-to-use telephone system only because that system does not need to understand what it is transmitting.

The same is true of the Internet, and of the internals of a PC. These technologies are designed to transmit and process information in any form, thereby providing tremendous cultural and economic value to their users. And the speed, power, and low cost of the Internet and PCs are possible precisely because they are designed to operate without having to understand the content of the data they are handling.

The general-purpose nature of the PC and the Internet is what has made them such astonishing engines of creativity, because it allows them to be used for purposes that their creators did not envision. Alexander Graham Bell did not foresee the invention of answering machines or voice mail – but he did not have to, because his general-purpose invention could accommodate them. He did not foresee that vending machines would phone a supplier when they ran out of candy bars. The designers of the Internet did not foresee the World Wide Web; but because the Internet infrastructure was general-purpose, the Web could rely on it immediately and without difficulty. General-purpose technologies provide platforms for innovation that allow anyone, even the proverbial kid in a garage, the opportunity to develop the next "killer app."

Copy protection operates on the opposite theory, by requiring the technology to categorize and understand the data that it is handling. Because copy protection and general-purpose computing and networking are fundamentally incompatible, attempts to add copy protection to general-purpose computers and networks are doomed to failure.

History supports this conclusion. To my knowledge, every copy protection scheme for general-purpose computers that has undergone serious public scrutiny has been found to be ineffective. For example, in the fall of 2000 the music industry, under the umbrella of the Secure Digital Music Initiative (SDMI), proposed a technological standard, based on watermarking, for copy protection of recorded music. When the SDMI challenged the public to analyze their technology, a team of independent researchers (including me) found that the technology could be defeated easily. The SDMI was wise to submit its technologies to public scrutiny, because this enabled the flaws in those technologies to be discovered and discussed before the technologies were widely deployed, thereby avoiding the high cost of investing in a doomed approach. The industry spent years of effort designing these technologies, but our small team of researchers was able to discredit the technologies thoroughly in less than three weeks. This was not because we were smarter than those technologies' designers, but simply because they were trying to do the impossible.

This general storyline has been repeated over and over through the years, as one PC copy protection technology after another has been abandoned after collapsing like a house of cards. The history of PC copy protection consists of cycles of overconfidence followed by disillusionment. Is the content industry overconfident about their current technology? They tell us that they will succeed, that their best people are working the problem. Others have said similar things in the past and have always turned out to be wrong.

Space does not permit me to categorize here the various approaches to PC copy protection, or to debunk them one by one. Many approaches are possible and might be proposed as a potential standard. My point is that whatever approach is proposed, there are fundamental difficulties, both theoretical and practical, that make PC copy protection very unlikely to work; and that in any case a proposed technology ought to undergo public scrutiny before anyone seeks to impose it on the entire information technology industry.

Of course, we cannot rule out the possibility that the advocates of PC copy protection, who have always been wrong in the past, will turn out to be right this time. Though unlikely, that is possible. But it is a mistake to blindly assume that they will turn out to be right, because having a broken standard is worse than having no standard at all.

Consider what will happen if a government-mandated protection measure turns out not to work. Such a measure would do many things: it would inconvenience honest consumers; it would raise the price of media players; it would lengthen product development cycles; it would impede the development of new and better standards. Everyone would suffer, except the pirates. The industry that devised the measure would look technically inept, and the government that mandated its use would look worse.

In an attempt to sweep all of this under the rug, the content industry has framed the issue cleverly as one of standardization. This presupposes that there is a menu of workable technologies, and the only issue is which of them to choose. They want us to ask which technology is best. But we should instead ask another question: Are any of these technologies workable in the first place? If not, then a standard for copy protection is as premature as a standard for teleportation.

In light of these facts, it would be a serious mistake to rush ahead and mandate a standard technology before that technology has been shown to provide any real protection for copyrighted material on PCs. The burden of proof should be on the advocates of PC copy protection to show that this time is different, that this time PC copy protection will work, for once. If a proposed standard can stand up to public scrutiny, then and only then will it be time to consider mandating it. If we mandate an unproven technology, then we are failing to learn from the unsuccessful history of PC copy protection – and we are indeed doomed to repeat it.

Contact information:
Edward W. Felten
Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

(650) 723-0366 (voice)
(650) 723-8440 (fax)
ed@felten.com