Trusted Computing, Peer-To-Peer Distribution, and the Economics of Pirated Entertainment

Stuart E. Schechter, Rachel A. Greenstadt, and Michael D. Smith
Harvard University
{*stuart,greenie,smith*}*@eecs.harvard.edu*

May 16, 2003

Abstract

The entertainment industry, facing a formidable threat from peer-to-peer piracy networks, is exploring every possible means to attack these networks. The industry is also employing defensive strategies to protect media and media players from those who would extract and copy their content. These content protection systems depend on the computer industry's newly announced 'trusted computing' technologies. While 'trusted computing' technologies may better protect media and media players from content extraction by pirates, we assert that the very same technologies can be employed to better protect pirates and their peer-to-peer distribution networks from the entertainment industry.

1 Introduction

The viability of content piracy hinges on the resource costs of and risk from two required steps: extracting content from its protected form and then distributing copies of that content. History demonstrates that advances in technology often reduce these costs. The latest such advance comes in the form of extraction tools and peer-to-peer networks that automate both steps of the piracy process and put them in the hands of the average consumer. In response, the entertainment industry is looking to protect their content using 'trusted computing' technologies, which aims to place content extraction technology back outside the reach of the average consumer. We explore the implications of such technologies and argue that history, against the hopes of the entertainment industry, may continue to repeat itself.

1.1 A brief economic history of piracy

The cost of pirated goods is a function of the costs of extracting content and distributing copies. We refer to the one-time extraction cost as e (sometimes

called the *first-copy* cost) and the per-copy distribution cost as d. The total per-copy cost of pirating n copies thus equals $\frac{e}{n} + d$, where the cost of extraction is amortized over the number of copies. Using this simple formula as a guide, we briefly review the evolution of the economics of piracy and set a framework for understanding the reasoning behind the anti-piracy techniques used in the past and those being proposed today.

Before the days of consumer-writable media, the cost of piracy was dominated by the per-copy distribution cost d. No effort was expended to make it costly to extract content from media. This one-sided approach makes sense when one considers the components of the distribution cost d: the resource costs related to purchasing and writing media and the legal liability costs associated with the distribution of pirated content in countries that enforce intellectual property laws. The direct effect of high resource costs is to limit the number of pirates. Because the average consumer could not afford to produce pirated media, the entertainment industry could easily afford to pursue legal action against those few with the financial resources for engaging in piracy. Such legal actions had the effect of increasing liability, which ultimately resulted in further increases in per-copy distribution costs.

The advent of audiotape and videotape made recording technology and media available at a reasonable cost, and the widespread acceptance of consumer VCRs created a demand for pirated video content.[1] These technology changes dramatically reduced d, and the entertainment industry reacted by endeavoring to increase e.

In particular, the industry introduced anti-piracy mechanisms into content-players and recorders in order to raise the cost of extraction high enough so that this cost could only be justified if amortized over a large number of copies. Consumer VCRs were built with technology that would refuse to record audio and video signals from sources of copyrighted content [8]. In parallel, the entertainment industry also employed patent protection and industry license agreements to force manufacturers to include anti-piracy mechanisms in their content players. These legal barriers were meant to exclude from the content-player market any manufacturer not complying with the anti-piracy design requirements. Increasing e made casual piracy prohibitively expensive, and the entertainment industry again kept piracy at bay by investigating and prosecuting only a small number of distributors.

The development of digital content players and cheap digital media again dramatically changed the economics of piracy by driving the resource costs related to purchasing and writing media to near zero.[2] In addition, digital media eliminated the problem of copy degradation and further drove down the costs of distribution. At first, the entertainment industry reacted by delaying the

[1] Even though the proliferation of pirated content was limited by imperfections introduced as copies of copies were made on analog media, these consumer technologies reduced d to the point where the number of potential pirates could increase dramatically.

[2] At the time of this writing, storage costs were approximately 30 cents per gigabyte for removable media, such as DVDs/CDs, and $1 per gigabyte for fixed storage, such as hard disks.

introduction of high-density, writable digital media into the consumer market. However, once personal computers (PCs) advanced to the point where compressed audio and video was easy to play and distribute across the Internet, it no longer made economic sense to block the sale of high-density, writable drives to consumers. Writable CD-ROM drives are now standard equipment on PCs, and drives that also write to DVD will soon take their place.

A primary goal of the DVD format was to protect digital video from piracy. As with VCRs, legal barriers and economic incentives were put in place to ensure that manufacturers could only produce a DVD reader if it included anti-piracy mechanisms to thwart content extraction and reverse engineering [2, page 431]. Once again, the industry's legal efforts would then focus on a smaller set of larger pirate distributors. For these reasons the industry has fiercely protected the DVD format, filing suit under the new Digital Millennium Copyright Act (DMCA) to keep video content extraction tools out of the hands of consumers [11, 12]. The entertainment industry has also tried, rather unsuccessfully, to retrofit the CD format with similar content-extraction protections [4].

Napster was the first system to integrate the end user into the distribution process. The reduction in the per-copy cost of pirated content was so significant that the market for pirated music and video content exploded. The market growth was aided by an image of legitimacy resulting from extensive press coverage and professional looking software. Having failed to protect content on CDs, the recording industry attacked the distribution channel, suing Napster as it would any other large distributor of pirated content. Though Napster's centralized infrastructure failed to survive legal attack, newer systems such as Gnutella and Kazaa evolved to use distributed infrastructures more resilient to legal action against individual components. While the Recording Industry Association of America (RIAA) is working to bring makers of piracy applications into US jurisdiction [5] and break the corporate veil [19], these piracy networks are designed to live on long after the demise of their creators.

Without an effective way to raise extraction costs or eliminate the current peer-to-peer distribution channels using legal attacks, the entertainment industry has undertaken a two-pronged effort to raise the per-copy distribution cost seen by individual consumers. On the legal front, the industry is using high profile litigation against a few individuals, in hopes of raising in all consumers the perceived liability of using these networks [18]. It is a strategy that appears to be having an effect [14]. The industry is also learning to use a technical approach to raising distribution costs. In particular, it is attacking the confidentiality, integrity, and availability of peer-to-peer distribution networks.

1.2 Enter 'trusted computing'

While attacking channels for distributing pirated content has not been without benefit, it also has costs and limitations. Thus, the entertainment industry continues to explore new ways of protecting the content stored on media and played by software. In particular, 'trusted computing' technologies promise to enable media players within a PC to execute with the same level of resistance

to piracy that one would expect from a proprietary hardware player, such as those used to play DVDs. If these technologies succeed, extracting content from the media of the future will be significantly more difficult than ripping a CD is today.

Part of the success of the entertainment industry's anti-piracy effort relies on its ability to make content extraction inconvenient enough to deter the general public. To be successful, the industry must also deter those individuals and defeat those systems that distribute pirated content. In short, the industry would like to return to the days when investigation and legal actions were sufficient to counter a reasonably sized set of professional pirates.

1.3 Roadmap

The per-copy cost of piracy, $\frac{e}{n} + d$, is at the heart of the ongoing battle between the entertainment industry and content pirates. In Section 2 we explain how 'trusted computing' technologies will be used to protect media players from content-extraction attacks, increasing the pirate's cost of extraction, e. We describe attacks that may be employed against peer-to-peer distribution of pirated content in Section 3. If successful, these attacks will increase the pirate's distribution costs, d, and reduce the number of copies, n, that the network is able to distribute. In Section 4, we explore a how the 'trusted computing' technologies described in Section 2 can be used by pirates to secure their peer-to-peer networks against the attacks of Section 3.

2 Protecting Content

To protect their content, owners will encrypt it before writing it to media or otherwise transmitting it to media players. Media players will be required to provide a minimum level of resistance to content-extraction attacks before content-owners will entrust them with the decryption keys. Because the PC platform was not designed to resist such attacks, media players running on today's PCs cannot make such guarantees. Not surprisingly, the leading forces in the PC market formed the Trusted Computing Platform Alliance (TCPA), now succeeded by the Trusted Computing Group (TCG), to introduce technologies that will enable PCs and their applications to obtain the trust of the entertainment industry. Microsoft has also introduced similar technologies as part of its next-generation secure computing base for Windows, formerly known as Palladium.

These efforts introduce into commodity computing hardware a private key of a public key pair, as described in Arbaugh, Farber, and Smith's early work on secure boot processes [3]. After placing the private key into the hardware, the manufacturer creates a signed certificate vouching that the hardware into which the key was placed exhibits certain properties, such as tamper-resistance, and that only this hardware was given the public key. The hardware may make claims, or *attest* to statements, to a remote entity by signing these claims with

it's private key. Trust in the claims certified by this *remote attestation* [1] process is only as strong as the trust in the entities that has signed off on the claims. Once claims regarding the identity and anti-piracy properties of the hardware and BIOS have been established, the BIOS may then attest to the identity of the code it will next execute, the operating system. In a final transitive step, an operating system trusted by the remote entity may then attest to the identity and integrity of the application it is running. In order to reduce the number of digital signatures required, hardware registers may be used to collapse these steps into a single claim by the hardware. Alternative approaches place full responsibility for protecting clients in the hardware, removing the need for attestation of the operating system [17].

If each link in the chain is trustworthy then a remote entity may rely upon a client application to behave with the trust properties, such as resistance to content-extraction, for which the application has been certified. Because operating systems rely upon hardware for their correct operation, and applications rely upon operating systems for their correct operation, each attestation step builds on the prior trust layers. If any layer turns out not to be trustworthy, it may subvert all the layers above it.

Once a trust infrastructure is in place, the entertainment industry may protect its content by encrypting it and only transmitting the keys to those platforms built from components (hardware, operating system, and applications) that it trusts. In order to ensure the confidentiality of the keys that protect content and the unencrypted content itself, additional operating services are required to protect them while applications use them. Specifically, the operating system must protect the applications's memory and, if keys are to be stored locally, its file storage. Operating system services will also be required to protect the content on its way to the screen or audio card, lest content be stolen in a digital format on its way to the user. Microsoft's next-generation secure computing base for Windows provides each of these services under the names *curtained memory*, *secure storage*, and *secure input and output*.

However, if humans are to eventually hear the protected audio signals and view the protected video signals, then this protected content can also be recorded. Since video cameras and music recorders can record and store any information perceivable to human eyes and ears, secure output paths all the way from computer to user are therefore impossible. A motivated attacker, who purchases the highest quality viewing or listening equipment and pairs it with equipment that can record the experience, will be able to produce a copy that is good enough to please a vast number of consumers. These limitations are acceptable if the goal is only to increase the cost of extraction enough to deter consumers, not professional pirates, from making copies.

3 Attacking Peer-to-Peer Distribution

Because no level of media protection can raise the cost of extraction beyond the cost of recording the signal presented to the user, a successful anti-piracy

effort must also work to maintain a high cost of distributing pirated content. In particular, the entertainment industry must determine how it can deter peer-to-peer distribution of its pirated content.

We explore attacks on peer-to-peer networks and the countermeasures used to defeat them. We consider these attacks with regard to the security assets they target: confidentiality, integrity, and availability.

3.1 Confidentiality

Breaches of confidentiality both increase the expected liability cost of distributing content and reveal information that can be used to write programs that attack the system's integrity and availability.

If caught, both senders and receivers of pirated content may face lawsuits or other forms of retaliatory action. Using today's peer-to-peer networks is particularly risky because anyone eavesdropping between the sender and the receiver may observe pirated content in transit. Even if content was transmitted in encrypted form, the eavesdropper could use traffic analysis to determine the network addresses of the sender and the receiver and the size of the files being transferred. These attackers use confidentiality attacks to interrupt file transfers [6], locate pirates in order to send them cease and desist messages [13], and gather evidence for litigation.

The first step in protecting the confidentiality of the network is to encrypt the data sent over it so that only the sender and receiver know what was sent. However, there is nothing encryption can do to ensure that the party at the other end of the line, who knows what was transmitted, is not the attacker. For this reason systems that provide anonymity, or at least plausible deniability, are desirable. In such systems, the attacker may know that copyrighted content was transmitted through the network but cannot identify the original sender or final recipient.

A common approach to anonymous networking is to re-route communications through more nodes than can be tracked effectively [20, 21]. Attackers may watch the communication as it travels through the network or run routers that expose routing information, but these threats may be mitigated so long as a reasonable fraction of the routers act to keep routing information confidential. At present, there is no way to determine which clients will route traffic through the network with the intent of protecting anonymity.

Attacking the network is not the only way to breach the confidentiality of the peer-to-peer system. By running the peer-to-peer client software and thus controlling a peer, an attacker may look into the peer-to-peer network through the "eyes" of its client software. Client software has no secrets because operating systems make every byte of a program's memory available to the machine's administrator, or root account. The attacker can locate encryption keys, network topology information, or any of the other information required to participate in the peer-to-peer network. Once confidentiality has been breached, the attacker may use the information to write programs to impersonate a genuine peer-to-peer client and attack the network from within. Such programs are invaluable

to the attacker as they enable scalable attacks on integrity and availability.

3.2 Integrity

The integrity of information in a peer-to-peer system may be attacked through the introduction of degraded-quality content or by misrepresenting the identity of the content. In the context of music, these attacks have included introducing noisy recordings or falsely labelling songs. Attacks on the integrity of information describing the operation of the peer-to-peer network, such as the network's topology and routing information, may disrupt communication or even prevent users from ever accessing the network again. If clients are disconnected from the network, or if content may be misrepresented or its quality decreased, then the user's cost of obtaining pirated content (part of the distribution cost) will increase.

Reputation systems counter corrupt content attacks by enabling users to rate the validity of content and those who provide it. To ensure that all copies of the same content share the same reputation, content may be identified by its fingerprint (or hash). This enables reputations to scale far beyond trust in the user and allows widely duplicated corrupt files to be recalled quickly.

To ensure that an attacker cannot modify or delete its client's reputation information, designers must distribute this information among the other clients using protocols that prevent tampering. Because attackers can delete clients and reinstall new ones, a reputation system should also maintain information for the machines on which clients run. Confounding this problem are virtual machines, in which the few potential unique machine identifiers (e.g. network card addresses) may be modified easily.

While we may construct reputation systems to be resilient to a large number of malicious users, no existing system is immune to attack from an unlimited number of such users [7, 16]. If the attacker can write programs that impersonate genuine clients, there is no limit to the number of malicious peers that can be introduced into the system.

3.3 Availability

More resources are expended performing searches on peer-to-peer networks than are required to request that a search be performed. Attackers may use their client application to issue a large number of search requests, flooding the network with more requests than can be serviced. Alternatively, the attacker may force their client application to drop packets it was meant to route by manipulating the operating system or by simply disconnecting network cables at the right times.

Peers can stem the flood of requests by requiring that requests be accompanied by proof that the requestor had performed computational work, restoring the balance between the computation costs of issuing and responding to requests. This approach was introduced by Dwork and Naor [10] to increase the low cost of sending email and make sending spam unprofitable. This concept

has been extended to more general settings, such as preventing network level denial of service attacks for TCP [15] and TLS [9]. Requiring clients to solve puzzles before issuing requests could go a long way to prevent flooding attacks on peer-to-peer networks. However, the entertainment industry might be able to harness enough processing power to flood networks if its members can exploit the media players they controls to perform puzzle computations on machines paid for by their users.

An alternative to client puzzles is to use the reputation systems mentioned above to track individual machine's utilization of networks resources. The efficacy of this approach is limited if the attacker can corrupt the reputation system using programs that impersonate genuine clients, or even if a large number of genuine clients can be run on virtual machines and fed scripted input. The payoff to the entertainment industry of scaling such attacks comes in the form of increased barriers between users and pirated content, which in turn increases the per-copy cost of distribution.

4 Defending Peer-to-Peer Distribution

At the time of this writing, Sharman Networks, the makers of Kazaa, claims that well over 200 million copies of its client application had been downloaded. Because these networks contain vast resources, attacks will only be affordable if the cost of attack is many times smaller than the damages inflicted on the distribution network.

The existing countermeasures described in Section 3 are sufficient to defend peer-to-peer networks against attacks from individual users running authentic clients on real machines. Attackers still have a leg up in that they may peer into clients running on their own machines, use this information to write programs that impersonate real clients, and run as many copies of these clients as they need to disrupt the network. Alternatively, they may script attack behaviors and feed those behaviors into a large number of authentic clients running in parallel on virtual machines.

Can peer-to-peer networks be made immune from malicious client software written by the attacker? They can if the personal computer industry delivers on its promise of remote attestation. Though this technology was envisioned to thwart pirates, it is exactly what a peer-to-peer system needs to ensure that no client application can enter the network unless that application, and the hardware (not a virtual machine) and operating system it is running on, has been certified by an authority trusted by the existing clients. The trust model may be quite simple: accept only new clients into the network if they are certified by the same authority that vouched for the existing clients.

What's more, if Microsoft delivers on the promises of its next-generation secure computing base for Windows, then clients can also be assured of secure storage and curtained memory. With these technologies, peer-to-peer systems can protect the confidentiality and integrity of the clients' memories, which are collectively the memory of the entire network.

5 Conclusion

To thwart piracy the entertainment industry must keep distribution costs high, reduce the size of distribution networks, and (if possible) raise the cost of extracting content. However, if 'trusted computing' mechanisms deliver on their promises, large peer-to-peer distribution networks will be more robust against attack and trading in pirated entertainment will become safer, more reliable, and thus cheaper. Since it will always be possible for some individuals to extract content from the media on which it is stored, future entertainment may be more vulnerable to piracy than before the introduction of 'trusted computing' technologies.

6 Acknowledgments

This paper could not have been completed without the advice, comments, and suggestions of Ross Anderson, Kim Hazelwood Cettei, Roger Dingledine, Glenn Holloway, David Molnar, Michael Rabin, and the anonymous reviewers. This research was supported in part by grants from Compaq, HP, IBM, Intel, and Microsoft.

References

[1] The Trusted Computing Platform Alliance. Building a foundation of trust in the PC. Technical report, January 2000.

[2] Ross J. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons, Inc., first edition, 2001.

[3] William A. Arbaugh, David J. Farber, and Jonathan M. Smith. A secure and reliable bootstrap architecture. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 4–7, 1997.

[4] John Borland. Customers put kibosh on anti-copy CD. *CNET News.Com*, November 19, 2002.

[5] John Borland. U.S. liability looms over Kazaa. *CNET News.Com*, November 25, 2002.

[6] John Borland. Fingerprinting P2P pirates. *CNET News.Com*, February 20, 2003.

[7] Fabrizio Cornelli, Ernesto Damiani, Sabrina De Capitani di Vimercati, Stefano Paraboschi, and Pierangela Samarati. Choosing reputable servents in a P2P network. In *Proceedings of The Eleventh International World Wide Web Conference*, May 7–11, 2002.

[8] Macrovision Corporation. Solutions > video technology > copy protection. http://www.macrovision.com/solutions/video/copyprotect/index.php3.

[9] Drew Dean and Adam Stubblefield. Using client puzzles to protect TLS. In *Proceedings of the 10th USENIX Security Symposium*, August 15–17, 2001.

[10] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Proceedings of Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference*, volume 740 of *Lecture Notes in Computer Science*. Springer, August 16–20, 1992.

[11] Evan Hansen. Ban on DVD-cracking code upheld. *CNET News.Com*, November 28, 2001.

[12] Amy Harmon. Judges weigh copyright suit on unlocking DVD shield. *The New York Times*, May 2, 2001.

[13] Amy Harmon. Music swappers get a message on PC screens: Stop it now. *The New York Times*, April 30, 2003.

[14] Amy Harmon. Record concerns sue to end piracy. *The New York Times*, April 23, 2003.

[15] Ari Juels and John Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *Proceedings of the 1999 Network and Distributed System Security Symposium*, February 4–5, 1999.

[16] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *Proceedings of The Twelfth International World Wide Web Conference*, May 20–24, 2003.

[17] David Lie, Chandramohan A. Thekkath, Mark Mitchell, Patrick Lincoln, Dan Boneh, John C. Mitchell, and Mark Horowitz. Architectural support for copy and tamper resistant software. In *ASPLOS-IX Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 168–177, November 12–15, 2000.

[18] Declan McCullagh. File-swapping foes exert P2P pressure. *CNET News.Com*, August 13, 2002.

[19] Stefanie Olsen. Record labels sue Napster investor. *CNET News.Com*, April 22, 2002.

[20] Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.

[21] Paul F. Syverson, David M. Goldschlag, and Michael G. Reed. Anonymous connections and onion routing. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 44–54, May 4–7, 1997.