
CHAPTER 6

Balance Toppled

Who Owns the Bits?

Automated Crimes—Automated Justice

Tanya Andersen was home having dinner with her eight-year-old daughter in December 2005 when they were interrupted by a knock at the door. It was a legal process server, armed with a lawsuit from the Recording Industry Association of America (RIAA), a trade organization representing half a dozen music publishers that together control over 90% of music distribution in the U.S. The RIAA claimed that the Oregon single mother surviving on disability payments owed them close to a million dollars for illegally downloading 1,200 tracks of gangsta rap and other copyrighted music.

Andersen's run-in with the RIAA had begun nine months previously with a "demand letter" from a Los Angeles law firm. The letter stated that "a number of record companies" had sued her for copyright infringement and that she could settle for \$4,000–\$5,000 or face the consequences. She suspected the letter was a scam, and protested to the RIAA that she had never downloaded any music. Andersen repeatedly offered to let the record companies verify this for themselves by inspecting her computer's hard drive, but the RIAA refused the offers. At one point, an RIAA representative admitted to her that he believed she was probably innocent. But, he warned, once the RIAA starts a lawsuit, they don't drop it, because doing so would encourage other people to defend themselves against the recording industry's claims.

Andersen found a lawyer after the December lawsuit was served, and they convinced a judge to order an inspection of the hard drive. The RIAA's own expert determined that Andersen's computer had never been used for illegal

downloading. But instead of dropping the suit, the RIAA increased the pressure on Andersen to settle. They demanded that their lawyers be allowed to take a deposition from Andersen's daughter, and even tried to reach the child directly by calling the apartment. An unknown woman phoned her elementary school principal falsely claiming to be her grandmother and asking about the girl's attendance. RIAA lawyers contacted Andersen's friends and relatives, telling them that Andersen was a thief who collected violent, racist

A great deal of information about digital copyright issues can be found at www.chillingeffects.org, a joint project of the Electronic Frontier Foundation and of several university law clinics.

music. The pressure on the 41-year-old Andersen, who suffered from a painful illness and emotional problems, forced her to abandon her hope of entering a back-to-work program. Instead, she sought additional psychiatric care. Finally, after two years, Andersen was able to file a motion for summary judgment, which

required the RIAA to come to court with proof of their claims. When they could not produce proof, the case was dismissed. Andersen is currently suing the RIAA for fraud and malicious prosecution.

26,000 Lawsuits in Five Years

The RIAA has filed more than 26,000 lawsuits against individuals for illegal downloading since 2003. The process begins when MediaSentry, RIAA's investigative company, logs into a file-sharing network in search of computers hosting music for download. MediaSentry connects to these computers and scans them for music files. When it finds something suspicious, it sends the computer's IP address to the RIAA's Anti-Piracy group, together with a list of the files it found. RIAA staff members download and listen to a few of these to verify that they are in fact copyrighted songs. Then they file a lawsuit against "John Doe," the person who uses the computer at the offending IP address. (See the Appendix for an explanation of IP addresses and other aspects of Internet structure.) With the lawsuit as a legal basis, they subpoena the computer's Internet Service Provider, forcing disclosure of the real name of the John Doe user at that IP address. The RIAA sends the user its demand letter, naming the songs that were verified and citing the total number of songs found as the basis for damages. The letter offers an opportunity to settle; the average settlement demand is about \$4,000, non-negotiable. There's even a web site, p2plawsuits.com, which users can visit to pay conveniently.

It's an automated sort of justice for the digital age. But these are automated sorts of crimes. File-sharing programs are commonly configured to start up and run automatically, exchanging files without human intervention.

The computer's owner may not even be aware that it has been configured to upload files in the background.

It's also an error-prone form of justice. Matching names to IP addresses is unreliable—several computers on the same wireless network might share the same IP address. An Internet Service Provider allocating IP addresses might shift them around, so that a computer with a particular IP address today might not be the same computer that was file sharing from that IP address last week. Even if it is the same computer, there's no way to prove who was using it at the time. And maybe there was a clerical error in reporting.

The RIAA knows that the process is flawed, but given their stake in stopping downloading, they see no choice. Not only are they seeing their products being distributed for free, but they themselves might be liable to lawsuits from artists for neglecting to protect the artists' copyrights. Explains Amy Weiss, RIAA Senior Vice President for Communications, "When you fish with a net, you sometimes are going to catch a few dolphin.... But we also realize that this cybershopping needs to stop." Besides Andersen, other snared "dolphin" included a Georgia family that didn't own a computer, a paralyzed stroke victim in Florida sued for files downloaded in Michigan, and an 83-year-old West Virginia woman who hated computers and who, as it turned out, was deceased.

The High Stakes for Infringement

Error or not, most people choose to pay when they get the demand letter. The cost of settling is less than the legal fees for contesting, and the cost of losing the lawsuit is staggering: damages of at least \$750 for each song downloaded. The 4,000-song contents of a 20GB iPod would be grounds for minimum damages of \$3 million—a thousand times the cost of purchasing those songs on iTunes. (A GB, or *gigabyte*, is about a billion bytes.)

Driftnet justice, automated policing of automated crimes, and three million dollars minimum damages for an iPod's worth of music are consequences of policies honed for

\$750 A SONG

The minimum damages that the court *must* award for infringement is \$750 per infringing act. In cases where the infringement can be shown to be "willful," damages could be as high as \$150,000 per infringement, or \$600 million for the 4,000 songs on an iPod. For defendants who can prove that they weren't even aware of the infringement, the court still *must* award at least \$200 per infringement—a "mere" \$800,000 for 4,000 songs.

a pre-networked world colliding with the exponentials of the digital explosion. Take the \$3-million iPod. This traces to the Copyright Act of 1976, which introduced a provision letting copyright holders sue for minimum *statutory damages* of \$750 per infringement.

The rationale for statutory damages is to ensure that the penalty is sufficient to deter infringement even when actual damages to the copyright holder are small. The scale of the damages has dreadful consequences in the age of digital reproduction, because each time a song is copied (uploaded or downloaded), it counts as a *separate* infringement. That way of reckoning “acts of infringement” may have seemed reasonable when the standards were set in pre-Internet 1976—when people could make only a few unauthorized copies, one by one. But the damage calculations balloon into unreality when a thousand songs can be downloaded to a home computer in a few hours over a high-speed network connection.

Although the digital explosion may have blown the legal penalties for infringement out of realistic proportion to the offense, it has also brought a more fundamental change: that the public is now concerned with copyright at all. Before the Internet, what could an ordinary person do to infringe copyright—make fifty photocopies of a book and sell them on the street corner? That would surely be infringement. But it would also be a lot of work, and the financial loss to the copyright holder would be insignificant.

Of all the dislocations of the digital explosion, the loss of the copyright balance is the most rancorous. Ordinary people can now effortlessly copy and distribute information on a massive scale. Listeners clash with a content industry whose economics relies on ordinary people not doing precisely that. As a

SENDING A MESSAGE

In October 2007, Jammie Thomas, a Minnesota single mother of two who earns \$36,000 a year, was found guilty of sharing 24 songs on the Kazaa file-sharing network ... and fined \$222,000: \$9,250 per song. This was the first of the RIAA's 16,000 lawsuits that went all the way to a jury trial. In the others, people settled or, as with Tanya Andersen, the case was dismissed or dropped. Given the legal statutory damages for infringement, Thomas's fine for 24 songs could have been anywhere between \$18,000 and \$3.6 million.

A juror interviewed afterward reported that there were people advocating for fines at both ends of that spectrum during deliberation: “We wanted to send a message that you don't do this, that you have been warned.”

Said the RIAA's lawyer after the verdict was read, “This is what can happen if you don't settle.”

result, millions of people are today vilified as “pirates” and “thieves,” while content providers are demonized as subverters of innovation and consumer freedom trying to protect their outdated business models.

Of all the dislocations of the digital explosion, the loss of the copyright balance is the most rancorous.

The war over copyright and the Internet has been escalating for more than 15 years. It is a spiral of more and more technology that makes it ever easier for more and more people to share more and more information. This explosion is countered by a legislative response that brings more and more acts within the scope of copyright enforcement, subject to punishments that grow ever more severe. Regulation tries to keep pace by banning technology, sometimes even before the technology exists. Single mothers facing mind-numbing lawsuits are merely collateral damage in that war today. If we cannot slow the arms race, tomorrow’s casualties may come to include the open Internet and dynamic of innovation that fuels the information revolution.

NET Act Makes Sharing a Crime

Copyright infringement was not even a criminal matter in the U.S. until the turn of the twentieth century, although an infringer could be sued for civil damages. Infringement with a profit motive first became a crime in 1897. The maximum punishment was then a year in prison and a \$1,000 fine. Things stayed that way until 1976, when Congress started enacting a series of laws that repeatedly increased the penalties, motivated largely by prompting from the RIAA and the MPAA (Motion Picture Association of America). By 1992, an infringement conviction could result in a ten-year prison sentence and stiff fines, but only if the infringement was done “for the purpose of commercial advantage or private financial gain.” Without a commercial motive, there was no crime.

That changed in 1994.

During the 1980s, MIT became one of the first universities to deploy large numbers of computer workstations connected to the Internet and open to anyone on campus. Even several years later, public clusters of networked powerful computers were not very common. In December 1993, some students in one of the clusters noticed a machine that was strangely unresponsive and was strenuously exercising its disk drive. When the computer staff examined this “bug,” they discovered that the machine was acting as a file-server bulletin board—a relay point where people around the Internet were

uploading and downloading files. Most of the files were computer games, and there was also some word-processing software.

MIT, like most universities, prefers to handle matters like this internally, but in this case there was a complication: The FBI had asked about this very same machine only a few days earlier. Federal agents had been investigating some crackers in Denmark who were trying to use MIT machines to break into National Weather Service computers. While measuring network traffic into and out of MIT, the Bureau had noticed a lot of activity coming from this particular machine. The bulletin board had nothing to do with the Denmark operation, but MIT felt that it had to tell the FBI what was happening. An agent staked out the machine and identified an MIT undergraduate, accusing him of operating the bulletin board.

The Justice Department seized on the case. The software industry was growing rapidly in 1994, and the Internet was just starting to enter the public eye—and here was the power of the Internet being turned to “piracy.” The Boston U.S. Attorney issued a statement claiming that the MIT bulletin board was responsible for more than a million dollars in monetary losses, adding “We need to respond to the culture that no one is hurt by these thefts and that there is nothing wrong with pirating software.”

What had occurred at MIT involved copyright infringement to be sure, but there was no commercial motive and hence no crime—no basis on which the Justice Department could act. There might have been grounds for a civil suit, but the companies whose software was involved were not interested in suing. Instead, the Boston U.S. Attorney’s office, after checking with their superiors in Washington, brought a charge of wire fraud against the student, on the grounds that his acts constituted interstate transmission of stolen property.

At the trial, Federal District Judge Stearns dismissed the case, citing a Supreme Court ruling that bootleg copies do not qualify as stolen property. Stearns chastised the student, describing his behavior as “heedlessly irresponsible.” The judge suggested that Congress could modify the copyright law to permit criminal prosecutions in cases like this if it so wished. But he emphasized that changing the rules should be up to Congress, not the courts. To accept the prosecution’s claim, he warned, would “serve to criminalize the conduct ... of the myriad of home computer users who succumb to the temptation to copy even a single software program for private use.” He cited Congressional testimony from the software industry that even the industry would not consider such an outcome desirable.

Two years later, Congress responded by passing the 1997 No Electronic Theft (NET) Act. Described by its supporters as “closing the loophole” demonstrated by the MIT bulletin board, NET criminalized any unauthorized copying with retail value over \$1000, commercially motivated or not. This

addressed Judge Stearns's suggestion, but it did not heed his caution: From now on, anyone making unauthorized copies at home, even a single copy of an expensive computer program, was risking a year in prison. After only two more years, Congress was back with the Digital Theft Deterrence and Copyright Damages Improvement Act of 1999. Its supporters argued that NET had been ineffective in stopping "piracy," and that penalties needed to be increased. The copyright arms race was in full swing.

The Peer-to-Peer Upheaval

The NET Act marked the first time that the Internet had triggered a significant expansion of liability for copyright infringement. It would hardly be the last.

In the summer of 1999, Sean Fanning, a student at Northeastern University, began distributing a new file-sharing program and joined his uncle in forming a company around it: Napster. Napster made it easy to share files, especially music tracks, over the Internet, and to share them on a scale never before seen.

Here is how the system worked: Suppose Napster user Mary wants to share her computer file copy of Sarah McLachlan's 1999 hit *Angel*. She tells the Napster service, which adds "Angel; Sarah McLachlan" to its directory, together with an ID for Mary's computer. Any other Napster user who would like to get a copy of *Angel*, say Beth, can query the Napster directory to learn that Mary has a copy. Beth's computer then connects directly to Mary's computer and downloads the song without any further involvement from the Napster service. The connecting and downloading are done transparently by Napster-supplied software running on Mary's and Beth's computers.

The key point is that previous file-sharing set-ups like the MIT bulletin board were so-called *centralized systems*. They collected files at a central computer for people to download. Napster, in contrast, maintained only a central directory showing where files on other computers could be found. The individual computers passed the files among themselves directly. This kind of system organization is called a *peer-to-peer* architecture.

Peer-to-peer architectures make vastly more efficient use of the network than centralized systems, as Figure 6.1 indicates. In a centralized system, if many users want to download files, they must all get the files from the central server, whose connection to the Internet would consequently become a bottleneck as the number of users grows. In a peer-to-peer system, the central server itself need communicate only a tiny amount of directory information, while the large network load for transmitting the files is distributed over




FIGURE 6.1 Underlying organization of traditional and peer-to-peer client-server network architectures. On the top, a traditional centralized file distribution architecture, in which files are downloaded to clients from a central server. On the bottom, a Napster-style peer-to-peer architecture in which the central server holds only directory information and the actual files are transmitted directly between clients without passing through the server.

the Internet connections of all the users. Even the slow connections common with personal computers in 1999 were enough for Napster's peer-to-peer system to let millions of users share music files ... which they did. By early 2001, two years after Napster appeared, there were more than 26 million registered Napster users. At some colleges, more than 80% of the on-campus network traffic could be traced to Napster. Students held Napster parties. You hooked up a computer to some speakers and to the Internet, invited your friends over—and for any song title requested, there it was. Someone among those millions of Napster users had the song available for downloading. This was the endless cornucopia of music; the universal jukebox.

The Specter of Secondary Liability

Universal though it may have been, this jukebox was collecting no quarters for the music industry. Previous escapades in file sharing, usually done on a small scale among friends, were barely annoyances from an economic perspective. Even the MIT bulletin board that engendered the No Electronic Theft Act had perhaps a few hundred users altogether. Napster was on a completely different scale, where anyone could readily share music files with a few hundred thousand “friends.” The recording industry recognized this immediately, and in December 1999, just a few months after Napster appeared, the RIAA sued it for more than \$100 million in damages.

Napster protested that it had no liability. After all, Napster itself wasn't copying any files. It was merely providing a directory service. How could you hold a company liable for simply publishing the locations of items on the Internet? Wasn't that publication just exercising freedom of speech? Unfortunately for Napster, the California Federal District Court didn't agree, and in July 2000, found Napster guilty of secondary copyright infringement (enabling others to infringe, and profiting from the infringement). A year later, after an unsuccessful appeal to the Ninth Circuit, the court ordered Napster's file-sharing service to shut down.

Napster was dead, but it had captured the imagination of the technical community as a striking demonstration of the power of the

SECONDARY INFRINGEMENT

Copyright law distinguishes between two kinds of secondary infringement. The first is *contributory infringement*—i.e., knowingly providing tools that enable others to infringe. The second is *vicarious infringement*—i.e., profiting from the infringement of others that one is in a position to control, and not preventing it. Napster was found guilty of both contributory and vicarious infringement.

Internet's fundamental architecture. No central machine controls the network; every machine in the network has equal rights to send any other machine a message. Machines connected to the Internet are, as the lingo has it, *peers*. The notion of the Internet as a network of peer machines communicating with each other directly—as opposed to a network of client machines mediated by central servers—was hardly new. Even the very first Internet technical specification, published in 1969, described the network architecture in terms of machines interacting as a network of peers. Systems incorporating peer-to-peer communication between larger computers had been in wide use since the early 1980s.

Napster showed that the same principle remained valid when the peers were millions of personal computers controlled by ordinary people. Napster's use of peer-to-peer was illegal, but it demonstrated the potential of the idea. Research and development in distributed computing took off. In 2000 and 2001, more than \$500 million was invested in companies building peer-to-peer applications. And transcending its roots as a technical network architecture, "P2P" became enshrined in techno-pop-culture-speak as a catchword for organizations of all types—including social, corporate, and political—that harness the power of myriad cooperating individuals without reliance on central authorities. As one 2001 review gushed, "P2P is a mindset, not a particular technology or industry."

Napster had also given an entire generation a taste of the Internet as universal jukebox for which people would clamor. Yet the recording companies, who worked together to combat illegal downloading, failed to collaborate to create a legal and profitable Internet music service to fill the vacuum left by Napster. Instead of capitalizing on file-sharing technology, they demonized it as a threat to their business. That technological rejectionism ratcheted up the rancor in the arms race, but it also did something even more short-sighted. The music companies surrendered a vast business opportunity to the profit of more imaginative entrepreneurs. Two years later, Apple would launch its iTunes music store, the first commercially successful music downloading service.

Sharing Goes Decentralized

In the meantime, new file-sharing schemes sprouted up that explored new technical architectures in attempts to tiptoe around liability for secondary infringement. Napster's legal Achilles's heel had been its central directory. As the court had ruled, control of the directory amounted to control of the file-sharing activity, and Napster was consequently liable for that activity. The new architectures got rid of central directories entirely. One of the simplest methods, called *flooding*, works like this: Each computer in the file-sharing network maintains a list of other computers in the network. When file-sharer

Beth wants to find a copy of *Angel*, her computer asks all the computers in its list. Each of those computers offers to send Beth a copy of *Angel* if it has one, and otherwise relays Beth's request to all the computers on *its* list, and so on, until the request eventually reaches a computer that has the file. Figure 6.2 illustrates the process. In contrast to the Napster-style architecture in Figure 6.1, there's no central directory. Distributed architectures like this are powerful because they can be extremely robust. The network will keep working even if many individual computers fail or go offline, as long as enough computers remain to propagate the requests.

CONTENT-DISTRIBUTION NETWORKS

The bare-bones flooding method sketched here is too simple to support practical large networks. But the success of decentralized peer-to-peer architectures has stimulated research into practical *content-distribution network* architectures that exploit the efficiency and robustness of peer-to-peer methods.




FIGURE 6.2 In contrast to Napster-style peer-to-peer systems illustrated earlier, decentralized file-sharing systems such as Grokster have no central directories.

No Safe Harbors

The companies building the new generation of file-sharing systems hoped these distributed architectures would also immunize them against liability for secondary copyright infringement. After all, once users had the software, what they did with it was beyond the companies' knowledge or control. So, how could the companies be held liable for what users did? To the recording industry, however, this was just Napster all over again: exploiting the Internet to promote copyright infringement on a massive scale. In October 2001, the RIAA sued the makers of three of the most popular systems—Grokster, Morpheus, and Kazaa—for damages of \$150,000 per infringement.

The three companies responded that they had no control over the users' actions. Moreover, their software was only one piece of the infrastructure that enabled file-sharing, and there were many other pieces. If the three software companies were liable, wouldn't makers of the other pieces be liable as well? What about Microsoft, whose operating system lets users of one computer copy files from other computers? What about Cisco, whose routers relay the unlicensed copyrighted material? What about the computer manufacturers, whose machines run the software? Wouldn't a ruling against the file-sharing network software companies expose the entire industry to liability?

The Supreme Court had provided guidance for navigating these waters with the landmark 1984 case *Sony v. Universal Studios*. In an episode that foreshadowed the *Grokster* suit 17 years later, the MPAA had sued Sony Corporation, charging Sony with secondary infringement for selling a device that was threatening to ruin the motion picture industry: the video cassette recorder. As the President of the MPAA thundered before Congress in 1982: "I say to you that the VCR is to the American film producer and the American public as the Boston strangler is to the woman home alone."

In a narrow 5 to 4 decision, the Supreme Court ruled in Sony's favor, holding that even though there was widespread infringement from people using VCRs

... the sale of copying equipment, like the sale of other articles of commerce, does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes. Indeed, it need merely be capable of substantial noninfringing uses.

The technology industries applauded. Here was a reasonably clear criterion they could rely on in evaluating the risk in bringing new products to market. Showing that a product was capable of substantial noninfringing uses would provide a "safe harbor" against allegations of secondary infringement.

This 1984 scenario—a new technology, a threatened business model—was now being replayed in the 2001 *Grokster* suit. The file-sharing companies were quick to cite the *Sony* ruling in their defense, explaining that there were many noninfringing uses of file sharing.

In April 2003, the Central California Federal District Court agreed that this case was different from *Napster*, and dismissed the suit, citing the *Sony* decision and commenting that the RIAA was asking the court to “expand existing copyright law beyond its well-drawn boundaries.” In reaction, the RIAA immediately began its campaign of suing individual users of the file-sharing software—the campaign that would later snag Tanya Andersen and Jammie Thomas.

The District Court’s ruling was appealed, and it was upheld by the Ninth Circuit, the same court that had ruled against *Napster* three years earlier:

In short, from the evidence presented, the district court quite correctly concluded that the software was capable of substantial noninfringing uses and, therefore, that the *Sony-Betamax* doctrine applied.

The RIAA naturally appealed, and when the Supreme Court agreed to review the decision, the entire networked world held its breath. Were content publishers to have no legal recourse against massive file-sharing? Would the *Sony* safe harbor be overturned? In June 2005, the Court returned a unanimous verdict in favor of the RIAA:

We hold that one who distributes a device with the object of promoting its use to infringe copyright, as shown by clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties.

A Question of Intent

The content industry had won, although it ended up with less than it had hoped for. The MPAA wanted the court to be explicit in weakening the *Sony* “substantial noninfringing use” standard. Instead, the court declared that the *Sony* case was not at issue here, and it would not revisit that standard. The file-sharing companies’ liability, the court said, stemmed not from the capabilities of the software, but from the companies’ intent in distributing it.

The technology industries (other than the three defendants, who were driven out of business) breathed an immediate sigh of relief that *Sony* had been left intact. But this was quickly followed by second thoughts. The

Grokster decision had opened up an entirely new set of grounds on which one could be held liable for secondary infringement. As the court ruled: “Nothing in *Sony* requires courts to ignore evidence of intent to promote infringement if such evidence exists.”

But what evidence? If someone accuses your company of secondary infringement, how confidently can you defend yourself against accusations of bad intent? The *Sony* safe harbor doesn’t seem so safe any more.

Take an example: The *Grokster* ruling cited “advertising an infringing use” as evidence of an active step taken to encourage infringement. Apple introduced the iTunes desktop with its CD-copying software in 2001. Early advertisements heavily promoted the product with the slogan “Rip, Mix, Burn.” Was that a demonstration of Apple’s bad intent? Many people certainly thought so, including the Chairman of Walt Disney when he told Congress in 2002, “There are computer companies, that their ads, full-page ads, billboards up and down San Francisco and L.A., that say—what do they say?—‘rip, mix, burn’ to kids to buy the computer.”

Can your company risk introducing a product with that slogan in the post-*Grokster* era? You might expect that you would have every chance of winning

NO COMMERCIAL SKIPPING

In 2001, ReplayTV Network introduced a digital video recorder for television programs that included the ability to skip commercials automatically. It also permitted people to move recorded shows from one ReplayTV machine to another. The company was sued for secondary infringement by the major movie studios and television networks, and driven into bankruptcy before the case was concluded. The company that bought Replay’s assets settled the case, promising not to include these features in its future models.

an “intent” fight in court, but the risks of losing are catastrophic. In personal infringement cases like Tanya Andersen’s, even the minimum statutory damage penalties of \$750 per infringement could have meant a million dollar claim over the (falsely alleged) songs on her hard drive—a staggering burden for an individual. But a technology company could conceivably be liable for damages based on *every* song illegally copied by *every* user of a device. Say you sell 14 million iPods (the number Apple sold in 2006) times 100 songs allegedly copied per iPod times \$750 per song. That’s more than a trillion dollars in damages—more than 100 times the *total* retail revenues of the recording

industry worldwide in 2006! Liability like that might seem ridiculous, but that’s the law. It means that guessing wrong is a bet-the-company mistake.

Better to be conservative and not introduce products with features that might prompt a lawsuit, even if you are reasonably sure that your products are legal.

We can speculate about products and features that are unavailable today due to the uncertainties in *Grokster's* “intent” standard, coupled with penalties for secondary infringement penalties that could lead to nightmarish fines. Companies are naturally reluctant to give examples, but one might ask why songs shared wirelessly with Microsoft Zune players self-destruct after three plays, or why Tivo recorders don't have automatic commercial skipping or let you move recorded movies to a PC. Non-coincidentally, in 2002, the CEO of a major cable network characterized skipping commercials while watching TV as theft, although he allowed that “I guess there could be a certain amount of tolerance for going to the bathroom.”

But speculating about the consequences of liability alone is largely pointless, because these liability risks have not been increasing in a vacuum. A second front has opened up in the copyright wars. Here, the weapons are not lawsuits, but technology.

Authorized Use Only

Computers process information by copying bits—between disk and memory, between memory and networks, from one part of memory to another. Actually, most computers are able to “keep” bits in memory only by recopying them over and over, thousands of times a second. (Ordinary computers use what is called Dynamic Random Access Memory, or DRAM. The copying is what makes it “dynamic.”) The relation of all this essential copying to the kind of copying governed by copyright law has been intellectual fodder for legal scholars—and for lawyers looking for new grounds on which to sue.

Computers cannot run programs stored on disk without copying the program code to memory. The copyright law explicitly permits this copying for the purpose of running the program. But suppose someone wants simply to *look at* the code in memory, not to run it. Does that require explicit permission from the copyright holder? In 1993, a U.S. Federal Circuit Court ruled that it does.

Going further, computers cannot display images on the screen without copying them to a special part of memory called a display buffer. Does this mean that, even if you purchase a computer graphic image, you can't view the image without explicit permission from the copyright holder each time? A 1995 report from the Department of Commerce argued that it does mean exactly this, and went on to imply that almost any use of a digital work involves making a copy and therefore requires explicit permission.

Digital Rights and Trusted Systems

Legal scholars can debate whether copyright law mandates a future of “authorized use only” for digital information. The answer may not matter much, because that future is coming to pass through the technologies of digital rights management and trusted systems.

The core idea is straightforward. If computers are making it easy to copy and distribute information without permission, then *change computers* so that copying or distributing without permission is difficult or impossible. This is not an easy change to make; perhaps it cannot be done at all without sacrificing the computer’s ability to function as a general-purpose device. But it’s a change that’s underway nonetheless.

Here is the issue: Suppose (fictitious) Fortress Publishers is in the business of selling content over the Web. They’d like the only people getting their content to be those whose pay. Fortress can start by restricting access on their web site to registered users only, by requiring passwords. Much web content is sold like this today—for instance, *Wall Street Digest* or *Safari Books Online*. The method works well (or at least has worked well so far) for this type of material, but there’s a problem with higher-value content. How does Fortress prevent people who’ve bought its material from copying and redistributing it?

One thing Fortress can do is to distribute their material in encrypted form, in such a way that it can be decrypted and processed only by programs that obey certain rules. For instance, if Fortress distributes PDF documents created with Adobe Acrobat, it can use Adobe LiveCycle Enterprise Suite to control whether people reading the PDF file with Adobe Reader are allowed to print it, modify it, or copy portions of it. Fortress can even arrange to make a document “phone home” over the Internet—i.e., to notify Fortress whenever it is opened and report the IP address of the computer that is opening it. Similarly, if Fortress prepares music files for use with Windows Media Player, it can use Microsoft Windows Media Rights Manager to limit the number of times the music can be played, to control whether it can be copied to a portable player or a CD, force it to expire after a certain period of time, or make it phone home for permission each time it’s played so that the Fortress web server can check a license and require payment if necessary.

The general technique of distributing content together with control information that restricts its use is called *digital rights management* (DRM). DRM systems are widely used today, and there are industry specifications (called *rights expression languages*) that detail a wide range of restrictions that can be imposed.

DRM might appear to solve Fortress’s problem, but the approach is far from airtight. How can Fortress be confident that people using their material are

using it with the intended programs, the ones that obey the DRM restrictions? Encrypting the files helps, but as explained in Chapter 5, attackers break that kind of encryption all the time—it happens regularly with PDF and Windows Media. More simply, someone could modify the document reader or the media player program to save unencrypted copies of the material as they are running, and then distribute those copies all over the Internet for anyone's use.

To prevent this, Fortress could rely on the computer operating system to require that any program manipulating their content must be certified. Before a program is run, the operating system checks a digital

signature for the program to verify that the program is approved and has not been altered. That's better, but a really clever attacker might alter the operating system so that it will run the modified program anyway. How could anyone prevent that? The answer is to build a chip into every computer that checks the operating system each time the machine is turned on. If the operating system has been modified, the computer will not boot. The chip should be tamper-proof so that any attempt to disable it will render the machine inoperable.

This basic technique was worked out during the 1980s and demonstrated in several research and advanced development projects, but only since 2006 has it been ready for wide deployment in consumer-grade computers. The required chip, called a *Trusted Platform Module* (TPM), was designed by the *Trusted Computing Group*, a consortium of hardware and software companies formed in 1999. More than half of the computers shipped worldwide today contain TPMs. Popular operating systems, including Microsoft Windows Vista and several versions of GNU/Linux, can use them for security applications. One application, *trusted boot*, prevents the computer from booting if the operating system has been modified (for example, by a virus). Another application, called *sealed storage*, lets you encrypt files in such a way that they can be decrypted only on particular computers that you specify. Given today's concerns over viruses and Internet security, it's a safe bet that TPMs will become pervasive. One industry estimate shows that more than 80% of laptop PCs will include TPMs by 2009.

ENCRYPTION AND DRM

Chapter 5 explains public-key encryption and digital signatures—the technologies that make public distribution of encrypted material possible. The "messages" that Alice and Bob are exchanging might be not text messages, but rather music, videos, illustrated documents, or anything at all. As the first koan says, "it's all just bits." Thus, the encryption technologies that Alice and Bob use for secret communication can be used by content suppliers to control the conditions under which consumers can watch movies or listen to songs.

Asserting Control Beyond the Bounds of Copyright

Fortress Publishers' problem could be solved in a world of digital rights management reinforced by trusted computing, but is that something we should welcome?

For one thing, it gives Fortress a level of control over use of its material that goes far beyond the bounds of copyright law. When we buy a book today, we take for granted that we have the right to read it whenever we like and as many times as we like; read it from cover to cover or skip around; lend it to a friend; resell it; copy out a paragraph for use in a book report; donate it to a school library; open it without "phoning home" to tell Fortress we are doing so. We need no permission to do any of these things. Are we willing to give up these rights when books are digital computer files? How about music? Videos? Software? Should we care?

Now leave to one side, for a moment, the dispute between music companies and listeners. DRM and trusted computing technologies, once standard in personal computers, will have other uses. The same methods that, in one country, prohibit people from playing unlicensed songs can, in another coun-

The same methods that, in one country, prohibit people from playing unlicensed songs can, in another country, prevent people from listening to unapproved political speeches or reading unapproved newspapers.

try, prevent people from listening to unapproved political speeches or reading unapproved newspapers. Developers of DRM and trusted platforms may be creating effective technologies to control the use of information, but no one has yet devised effective methods to circumscribe the limits of that control. As one security researcher warned: "Trusted computing" means that "third parties can trust that your computer will disobey your wishes."

Another concern with DRM is that it increases opportunities for technology lock-in and anticompetitive mischief. It is tempting to design operating systems that run only certified applications in order to protect against viruses or bogus document readers and media players. But this can easily turn into an environment where no one can market a new media player without publishers' approval, or where no one can deploy *any* application without first having it registered and approved by Microsoft, HP, or IBM. A software company that poses a competitive threat to established interests, like publishers, operating system vendors, or computer manufacturers, might suddenly encounter "complications" in getting its products certified. One reason innovation has been so rapid in information technology is that the infrastructure is open: You don't need permission to

introduce new programs and devices on the Internet. A world of trusted systems could easily jeopardize this.

A third DRM difficulty is that, in the name of security and virus protection, we could easily slip into an unwinnable arms race of increasing technology lock-down that provides no real gain for content owners. As soon as attackers anywhere bypass the DRM to produce an unencrypted copy, they can distribute it—and they might be willing to go to a lot of effort to be able to do that.

Think, for example, about making unauthorized copies of movies. Very sophisticated attackers might modify the TPM hardware on their computers, putting a lot of effort into bypassing the tamper-proof chip. Here's an even easier method: let the TPM system operate normally, but hook up a video recorder in place of the computer display. That particular attack has been anticipated by the industry with a standard that requires all high-definition video to be transmitted between devices in encrypted form. Windows Vista implements this in its *Output Protection Management* subsystem, out of concern that otherwise the movie studios would not permit high-definition video to be played on PCs at all. Even that protection scheme is vulnerable—you could simply point a video recorder at the screen. The result would not be high-definition quality, but once it has been digitized, it could be sent around the Internet without any further degradation.

Content owners worried about these sorts of attacks refer to them as the *analog hole*, and there seems to be no technological way to prevent them. J.K. Rowling tried to prevent unauthorized Internet copies of *Harry Potter and the Deathly Hallows* by not releasing an electronic version of the book at all. That did not stop the zealous fan mentioned in Chapter 2 from simply photographing every page and posting the entire book on the Web even before it was in bookstores.

In the words of one computer security expert, “Digital files cannot be made uncopyable, any more than water can be made not wet.” There is one thing for certain: The DRM approach to copyright control is difficult, frustrating, and potentially fraught with unintended consequences. Out of that frustration has emerged a third response—along with liability and DRM—to the increasing levels of copying on the Internet: outright criminalization of technology.

Toward De-Escalation

Getting off that path requires freeing ourselves of old ideas and perspectives. Difficult as that seems, there are grounds for optimism. During 2007, the recording industry made a major shift away from reliance on digital rights

management. In addition to restraints it imposes on technology, DRM is an inconvenience both for consumers and publishers. There has been an increasing public acknowledgement of the downsides of DRM, not only by consumer groups, but by the industry itself.

One of the first visible moves was an announcement in February 2007 by Apple's Steve Jobs, in the form of an open letter to recording industry executives asking them to relax the licensing restrictions that required Apple to implement DRM on iTunes music. In Jobs's view, a world of online stores selling DRM-free music that could play on any player would be "clearly the best alternative for consumers, and Apple would embrace it in a heartbeat." The industry reacted coldly, but other groups chimed in to agree with Jobs. In March, Musicload, one of Europe's largest online music retailers, came out against DRM, noting that 75% of its customer service calls were due to DRM. Musicload asserted that DRM makes using music difficult for consumers and hinders the development of a mass market for legal downloads. In November, the British Entertainment Retailers Association also came out against DRM. Its director general claimed that copy protection mechanisms were "stifling growth and working against the consumer interest."

By the summer of 2007, Apple iTunes and (separately) Universal Music Group began releasing music tracks that could be freely copied. The iTunes tracks contained information ("watermarks") identifying the original purchaser from iTunes. That way, if large numbers of unauthorized copies would appear on the Internet, the original purchaser could be traced and held accountable.

A few months later, even that level of restriction was vanishing. By the beginning of 2008, all four major music labels—Universal, EMI, Warner, and Sony/BMG—were releasing music for sale through Amazon without watermarks that identified individual buyers. It was a remarkable about-face over the

course of a year. When Jobs made his February 2007 proposal, Warner Music CEO Edgar Bronfman flat-out rejected the idea as "completely without logic or merit." Before the end of the year, Warner was announcing that it would

USING WATERMARKING

Using watermarking rather than copy restrictions and access control is an example of a general approach to regulation through *accountability*, rather than *restriction*. Don't try to prohibit violations in advance, but make it possible to identify violations when they occur and deal with them then. The same perspective can apply in privacy, as mentioned in Chapter 2, where one can focus on the appropriate use of personal information rather than restricting access to it.

sell DRM-free music on Amazon, with Bronfman explaining in a note to employees:

By removing a barrier to the sale and enjoyment of audio downloads, we bring an energy-sapping debate to a close and allow ourselves to refocus on opportunities and products that will benefit not only WMG, but our artists and our consumers as well.

The increasing recognition that the DRM approach is failing is sparking experiments with other models for distributing music on the Internet. Universal has been talking to Sony and other labels about a subscription service, where users would pay a fixed fee and then get as much music as they want. One plan links the service to a new hardware device, here the price of the service would be folded into the price of the hardware.

A related idea is to distribute music through blanket licenses with mobile carriers or ISPs. New companies are emerging that offer this kind of service on college computer networks. Another variant is the idea of *unlimited content networks*. These are networks that give access to music or video that floats around the network with no restrictions. People can make unlimited use of the material—downloading, copying, moving it to portable devices, sharing with others—as long as they keep it within the network.

A complementary approach promotes sharing of music and other creative works in a way that enriches the common culture, by making it easy for creators to distribute their own work and to build on each other's work. One organization that provides technical and legal tools to encourage this is *Creative Commons*. This organization distributes a family of copyright licenses that creators can use for publishing their works on the Internet, including licenses that permit open sharing. The licenses are expressed both as legal documents and as computer code that can support new applications. If a work appears on the Web with the appropriate Creative Commons code, for example, search engines might return references to it when asked to find material that can be used under specified licensing conditions. Stimulating open sharing on the Internet is an example of moving toward a *commons*—that is, a system of sharing that minimizes the need for fine-grained property restrictions (Chapter 8, “Bits in the Air” includes more on the notion of a commons).

Experience with these and other approaches will show whether there are economically viable models for distributing music that do not rely upon

DRM. Success could pave the way for the motion picture industry and other publishers to get off the anti-circumvention path—a dead end that has been more effective at harming innovation than at stopping infringement, and which even some of the original architects of the policy are now acknowledging as a failed approach.

Even then, however, the larger problems created by the DMCA would not fade away, since policies locked into law are not easily unlocked. If the content industry moves to better business models and the DRM battles subside, the DMCA's anticircumvention provisions may continue to be anti-consumer, anti-competitive blots on the digital landscape. Unless repealed from the legal code, they would remain as battlefield relics of a war that was settled by peaceful means—unexploded ordnance that a litigious business could still use in ways unrelated to the law's original intent.

CREATIVE COMMONS LICENSES

If you've created works that you want to publish on the Internet, you can use the Creative Commons license generator at creativecommons.org to obtain a license tailored to your needs. With the license, you can retain specified rights of your choice while granting blanket permission for other uses.