EXHIBIT 4

IAAL*:
PEER-TO-PEER FILE SHARING AND COPYRIGHT LAW AFTER NAPSTER

by Fred von Lohmann
Senior Staff Attorney (Fair Use & Intellectual Property), EFF
fred@eff.org

What this is, and who should read it.

As the Napster saga illustrates, the future of peer-to-peer file-sharing is entwined, for better or worse, with copyright law. The legal fight has already broken out, with copyright owners targeting not only the makers of file-sharing clients like Napster and Scour, but also companies that provide products that rely on or add value to public P2P networks, such as MP3Board.com, which provides a web-based search interface for the Gnutella network.

The fight has only just begun. If these early skirmishes yield any lesson for future P2P developers, it's that a legal strategy needs to be in place early, preferably at the beginning of development, rather than bolted on at the end. As a result, if you are interested in peer-to-peer file sharing, whether as a developer, investor, or provider of ancillary services (such as search services, platform tools, or security), it's time to bone up on some copyright law basics.

This piece is meant as a general explanation of the U.S. copyright law principles most relevant to P2P file-sharing technologies. It is aimed primarily at:

> Developers of core P2P file-sharing technology, such as the underlying protocols, platform tools, and specific client implementations;
> Developers of ancillary services that depend upon or add value to P2P file-sharing networks, such as providers of search, security, metadata aggregation, and other services;
> Investors seeking to evaluate the potential copyright risks associated with the various ventures listed above.

The following discussion is meant as a general introduction, and thus occasionally glosses over some of copyright law's more subtle nuances. At the most basic level, it is aimed not at giving you all the answers, but rather at allowing you to recognize the right questions to ask your lawyers.

What this is not: The following discussion focuses only on U.S. copyright law, and does not address any issues that may arise under non-U.S. law. While non-copyright principles may also be mentioned, this discussion does not attempt to examine other legal principles that might apply to P2P file-sharing, including patent, trademark,

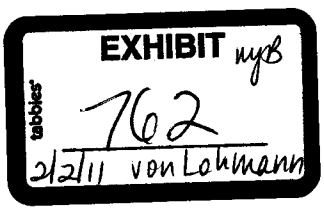

trade secret, or unfair competition. Nothing contained herein constitutes legal advice—please discuss your individual situation with your own attorney.

Copyright Basics and the Intersection with P2P Filesharing

Copyright law applies to virtually every form of expression that can be captured (or, to use the copyright term of art, "fixed") in a tangible medium, such as on paper, film, magnetic tape, hard drive, optical media, or even merely in RAM. Songs, books, photographs, software, and movies are all familiar examples of copyrighted works. Copyright protection begins from the moment that the expression is fixed, and continues for the lifetime of the author, plus 70 years.

During this period, copyright law reserves certain rights exclusively to the owner of the work, including the right to reproduce, distribute, and publicly perform the work. So, for example, if you wrote a song and recorded it on your computer, you would own the resulting copyrighted work and only you would have the right to make copies of the file, distribute it to the public, or sing the song in your local concert hall. If anyone else did any of these things without your permission, she would be infringing your copyright (unless the activity qualified as a "fair use" or fell into one of the other statutory exceptions to a copyright owner's exclusive rights).

The nature of digital file-sharing technology inevitably implicates copyright law. First, since every digital file is "fixed" for purposes of copyright law (whether on a hard drive, CD, or merely in RAM), the files being shared generally qualify as copyrighted works. Second, the transmission of a file from one person to another results in a reproduction, a distribution, and possibly a public performance (in the world of copyright law, "public performance" includes the act of transmitting a copyrighted work to the public). To a copyright lawyer, every reproduction, distribution, and public performance requires an explanation, and thus file-sharing systems seem suspicious from the outset.
The end-users: "direct" infringement.

For the individuals who are sharing files, the question becomes whether all of these reproductions, distributions, and public performances are authorized by the copyright owner or otherwise permitted under copyright law (as "fair use," for example). So, if the files you are sharing with your friends are videos of your vacation, you are the copyright owner and have presumably authorized the reproduction, distribution, and performance of the videos. However, if you are sharing MP3's of Metallica's greatest hits, or disc images of the latest Microsoft Office 2000 installation CD, the issue becomes more complicated. In that case, assuming that the copyright owner has not authorized the activity, the question of copyright infringement will depend whether you can qualify for any of the limited exceptions to the copyright owner's exclusive rights. If not, you're what copyright

lawyers call a "direct infringer"—you have directly violated one or more of the copyright owner's exclusive rights.

The P2P tool maker: "contributory" and "vicarious" infringement.

But what does this have to do with those who develop and distribute peer-to-peer file-sharing tools? After all, in a pure peer-to-peer file-sharing system, the vendor of the file-sharing tool has no involvement in the copying or transmission of the files being shared. These activities are handled directly between end-users. Copyright law, however, can sometimes reach beyond the direct infringer to those who were only indirectly involved in the infringing activity. As in many other areas of the law (think of the "wheel man" in a stick up, or supplying a gun to someone you know is going to commit a crime), copyright law will sometimes hold one individual accountable for the actions of another. So, for example, if a swapmeet owner rents space to a vendor with the knowledge that the vendor sells counterfeit CDs, the swapmeet owner can be held liable for infringement alongside the vendor.

Under copyright law, this indirect, or "secondary," liability can take two distinct forms: contributory infringement and vicarious infringement. In order to prevail under either theory, the copyright owner must first show that some underlying direct infringement has taken place. In other words, there must be a direct infringer before anyone will be "indirectly" liable. In a widely-used public peer-to-peer file-sharing environment, however, it is a virtual certainty that at least some end-users are engaged in infringing activity (unless specific technical measures are taken to prevent this, like permitting only the sharing of files that have been cryptographically marked as "authorized"). When the major record labels and music publishers decided to sue Napster, for example, it was not difficult for them to locate a large number of Napster users who were sharing copyrighted music without authorization.

Contributory Infringement

Contributory infringement is similar to "aiding and abetting" liability: one who knowingly contributes to another's infringement may be held accountable. Or, as the courts have put it, "one who, with knowledge of the infringing activity, induces, causes, or materially contributes to the infringing conduct of another, may be held liable as a contributory infringer." So, in order to prevail on a contributory infringement theory, a copyright owner must prove each of the following elements:

1. Direct Infringement: There has been a direct infringement by someone.
2. Knowledge: The accused contributory infringer knew of the underlying direct infringement. This element can be satisfied by showing either that the contributory infringer actually knew about the infringing activity, or that he reasonably should have known given all the facts and circumstances. At a minimum, however, the contributory infringer must have some specific

information about infringing activity—the mere fact that the system is capable of being used for infringement, by itself, is not enough.
3. Material Contribution: The accused contributory infringer induced, caused, or materially contributed to the underlying direct infringement. Merely providing the "site and facilities" that make the direct infringement possible can be enough.

Vicarious Infringement

Vicarious infringement is derived from the same legal principle that holds an employer responsible for the actions of its employees. A person will be liable for vicarious infringement if he has the right and ability to supervise the infringing activity and also has a direct financial interest in such activities. Thus, in order to prevail on a vicarious infringement theory, a copyright owner mush prove each of the following:

1. Direct Infringement: There has been a direct infringement by someone.
2. Right and Ability to Control: The accused vicarious infringer had the right and ability to control or supervise the underlying direct infringement. This element does not set a high hurdle. For example, the Napster court found that the ability to terminate user accounts or block user access to the system was enough to constitute "control."
3. Direct Financial Benefit: The accused vicarious infringer derived a "direct financial benefit" from the underlying direct infringement. In applying this rule, however, the courts have not insisted that the benefit be especially "direct" or "financial"—almost any benefit seems to be enough. For example, the Napster court found that "financial benefit exists where the availability of infringing material acts as a draw for customers" and the growing user base, in turn, makes the company more attractive to investors.

It should be noted that the nature of vicarious infringement liability creates a strong incentive to monitor the conduct of your users. This stems from the fact that knowledge is not required for vicarious infringement liability; a person can be a vicarious infringer even if they are completely unaware of infringing activity. As a result, if you exercise control over your users and derive a benefit from their activities, you remain ignorant of their conduct at your own risk. In the words of the Napster court, "the right to police must be exercised to the fullest extent. Turning a blind eye to detectable acts of infringement for the sake of profit gives rise to liability."

Indirect Liability and P2P Systems: the Napster Case

The Napster case represents the first application of these indirect liability theories to a peer-to-peer file-sharing service. In that case, the plaintiffs admitted that Napster did not, itself, make or distribute any or their copyrighted works. Instead, they

argued that Napster is liable for contributory and vicarious infringement. In its February 12, 2001 opinion, the Ninth Circuit agreed, rejecting each of Napster's proposed defenses.

Turning first to contributory infringement, the Ninth Circuit upheld the lower court's findings:

1. Direct Infringement: At least some Napster users are direct infringers, because they distributed and reproduced copyrighted music without authorization.
2. Knowledge: Napster had actual knowledge of infringing activity, based on internal company emails and the list of 12,000 infringing files provided by the RIAA. Moreover, Napster should have known of the infringing activity, based on the recording industry experience and downloading habits of its executives and the appearance of well-known song titles in certain promotional screen shots used by Napster.
3. Material Contribution: Napster provided the "site and facilities" for the directly infringing conduct of its users.

The Ninth Circuit also endorsed the lower court's vicarious infringement analysis:

1. Direct Infringement: At least some Napster users are direct infringers, because they distributed and reproduced copyrighted music without authorization.
2. Right and Ability to Control: Napster has the ability to control the infringing activity of its users because it retains the right to block a user's ability to access its system.
3. Financial Benefit: Napster derived a financial benefit from the infringing activities of its users because this activity acted as a "draw" for customers, and a portion of Napster's value is derived from the size of its user base.

The Ninth Circuit concluded, however, that the lower court had not adequately considered the technological limits of the Napster system when crafting the preliminary injunction. In ordering the district court to revise its injunction, the Ninth Circuit spelled out some guiding principles. First, in order to prevent contributory infringement, after receiving notice from a copyright owner that a work is being shared on its system without authorization, Napster will have to take reasonable steps to prevent further distribution of the work. Although the particulars will be up to the lower court, this almost certainly will require that Napster implement file-name filtering to its central index. It may also require that Napster implement more sophisticated filtering based on MP3 ID tags, MD5 hashes, acoustic fingerprints, or other meta-data.

Second, in order to prevent vicarious infringement, the Ninth Circuit declared that "Napster...bears the burden of policing its system within the limits of the system."

Again, the particulars of this command will be determined by the lower court. Nevertheless, this will almost certainly require some pro-active monitoring activity by Napster. Since, in the court's view, Napster "controls" its users, Napster will likely be required to take reasonable measures to keep tabs on what those users are up to, within the bounds of its system architecture. At a minimum, this will require that Napster pro-actively monitor its central index to weed out any songs that it knows are not authorized for sharing. It will also require that Napster continue to terminate users who share copyrighted works without authorization.

Potential Defenses Against Contributory and Vicarious Liability

No Direct Infringer: "All of My Users are Innocent Fair Users"

As discussed above, if there is no direct infringement, there can be no indirect liability. Consequently, if a peer-to-peer developer can plausibly claim that no users in the network are sharing copyrighted works without authorization, this would be a complete defense to any contributory or vicarious infringement claims. Unfortunately, this may be extremely difficult to demonstrate, given the decentralized nature of most P2P networks and the wide variety of uses to which they may be put. Even if file sharing by some users is privileged under the "fair use" doctrine or another statutory exception to copyright, it will be very difficult to show that every user falls within such an exception. Nevertheless, in certain specialized networks that permit the sharing of only secure, authorized file types, this may be a viable defense.

"Capable of substantial noninfringing uses"

Although contributory and vicarious infringement can sweep broadly, catching even those only indirectly involved in the copyright infringement of others, the Supreme Court has defined an outer limit to this reach. In a case involving the Sony Betamax VCR, the Supreme Court found that contributory infringement liability could not reach the manufacturer of a device that is "capable of substantial noninfringing use."

Universal City Studios and Disney were on one side of this case, arguing that the Sony Betamax VCR was a tool of copyright infringement. On the other side were Sony, its advertising agent, several of its retail distributors, and one individual VCR user. The case ultimately made its way to the Supreme Court, which ultimately issued a 5-4 decision that proceeded in two parts. First, the Court held that there could be no contributory liability, even if some VCR users were engaged in copyright infringement, so long as the device was "capable of a substantial noninfringing use." In the second part of its opinion, the Court found that the VCR was capable of several such noninfringing uses, including the time-shifting of television broadcasts by home viewers.

Unfortunately, the recent Ninth Circuit decision in the Napster case has dramatically reduced the scope of the "Betamax defense." First, the Napster court found that this defense does not apply to vicarious infringement liability. Accordingly, if you have control over, and derive a financial benefit from, direct infringement, the existence of "substantial noninfringing uses" for your service is irrelevant. Second, the court concluded that the Betamax defense only applies until specific information identifying infringing activity has been received. At that point, a failure to act to prevent the infringing activity will give rise to liability, and the existence of "substantial noninfringing uses" becomes irrelevant.

The Ninth Circuit's interpretation of the Betamax defense has at least two important implications for P2P developers. First, it underscores the threat of vicarious infringement liability—at least in the Ninth Circuit, a court will not be interested in hearing about your "substantial noninfringing uses" if you are accused of vicarious infringement. Accordingly, "control" and "direct financial benefit," as described above, should be given a wide berth. This will likely reduce the attractiveness of business models built on an on-going "service" or "community-building" model, to the extent that these models allow the provider to control user activity (i.e., terminate or block users) and create value by attracting a large user base. At the same time, it may increase the attractiveness of selling completely decentralized file-sharing software, insofar as this might minimize the vendor's "control" over, and continuing "direct financial benefit" from, infringing uses.

Second, with respect to contributory infringement, the Ninth Circuit's interpretation of the Betamax defense makes it risky to ignore "cease and desist" letters from copyright owners, which in turn may limit a developer's freedom to define the architecture of her product or service. Once you have received notice of specific infringing activity, your "substantial noninfringing uses" may no longer serve as a shield to contributory liability. Of course, even the Ninth Circuit recognized that the ability to respond to these notices may be limited by the technology behind the challenged service or product. Nevertheless, when a court is required to determine the limits of what is technically reasonable, the results can be uncertain. The Napster decision certainly suggests that copyright owners, once they make out a case of contributory or vicarious infringement liability, are in a position to demand that a developer of P2P tools take steps to reduce the likelihood that it will be used for infringing activity. What these steps might entail is difficult to predict, but may include, in some cases, modification of the architecture and capabilities of the tool, service or system.

The DMCA Section 512 "safe harbors"

In 1998, responding in part to the concerns of ISPs regarding their potential liability for the copyright infringement of their users, Congress enacted a number of narrow "safe harbors" for copyright liability. These safe harbors appear in section 512 of the Copyright Act, which in turn appears in title 17 of the U.S. Code (17 U.S.C. 512).

These safe harbors apply only to "online service providers," and only to the extent that the infringement involves four functions: transitory network transmissions, caching, storage of materials on behalf of users (e.g., web hosting, remote file storage), and the provision of information location tools (e.g., providing links, directories, search engines).

Each of these functions, however, is narrowly defined by the statute (e.g., they don't cover what you'd think) and reflects the state of the art in 1998. For example, the automated web page caching conducted by AOL in 1998 falls within the caching safe harbor, while the more sophisticated efforts of Akamai today may not. Because Congress did not anticipate peer-to-peer file sharing when it enacted the safe harbors, many P2P products may not fit within the four enumerated functions. For example, according to an early ruling by the district court in the Napster case, an OSP cannot use the "transitory network transmission" safe harbor unless the traffic in question passes through its own private network. Many P2P products will, by their very nature, flunk this requirement, just as Napster did.

In addition to being limited to certain narrowly-circumscribed functions, the safe harbors are only available to entities that comply with a number of complex, interlocking statutory requirements:

> The online service provider ("OSP") must (1) adopt, reasonably implement, and notify its users of a policy of terminating the accounts of subscribers who are repeat infringers; and (2) accommodate and not interfere with "standard technical measures" that have been widely adopted on the basis of industry-wide consensus (e.g., the use of robot.txt exclusion headers to block spiders).
> The OSP must designate a "copyright agent" to receive notices of alleged copyright infringement, register the agent with the Copyright Office, and place relevant contact information for the agent on its web site.
> The OSP must, upon receiving a notification of infringement from a copyright owner, expeditiously remove or disable access to the infringing material ("notice and takedown").
> The OSP must not have known about the infringement, or been aware of facts from which such activity was apparent (i.e., if you take a "head in the sand" approach, you lose the safe harbor).
> The OSP must not receive a direct financial benefit from infringing activity, in a situation where the OSP controls such activity (i.e., if you're liable for vicarious liability, the safe harbor may not protect you).

In the final analysis, qualifying for any of the DMCA safe harbors requires careful advance attention to the legal and technical requirements and obligations that the statute imposes. As a result, any P2P developer who intends to rely on them should seek competent legal counsel at an early stage of the development process—an after-the-fact, "bolt on" effort to comply is likely to fail (as it did for Napster). For more

detailed information regarding the limits and requirements of the safe harbors, you might consult the overview located at http://www.richmond.edu/jolt/v6i4/article1.html.

The DMCA ban on circumvention technologies

One recent addition to the copyright landscape deserves special attention. Section 1201 of the Copyright Act, enacted as part of the DMCA, makes it unlawful to "circumvent" any technology aimed at protecting a copyrighted work. In addition, the development, distribution or use of circumvention technology or devices is, with only narrow exceptions, also unlawful. For example, if a copyright owner uses a digital rights management ("DRM") solution to protect a song, it would be unlawful for anyone to crack the encrypted file without the copyright owner's permission, or to build or distribute a software tool designed to crack the file. The litigation involving DeCSS software, which is capable of decrypting video DVDs, represents one of the first cases testing these "anti-circumvention" provisions of the DMCA.

Of course, circumvention technology is not a necessary part of a peer-to-peer file-sharing network. Today's P2P protocols, such as Gnutella, simply facilitate file transfers, leaving the file itself, whether encrypted or not, unaltered. Nevertheless, as copyright owners begin to deploy DRM and watermarking systems, there may be interest in integrating circumvention tools with file-sharing tools. In light of the DMCA's broad ban on circumvention technology, however, any such integration may substantially increase the risk of liability.

Lessons and Guidelines for P2P Developers

A few general guidelines for P2P developers can be derived from the discussion above. These are steps you can take that may: (1) reduce the chance that your project will be an easy, inviting target for copyright owners; (2) placate your investors when they ask you whether you are likely to spend their money on litigation rather than products; and (3) minimize the chances that your case will become the next legal precedent that content owners can use to threaten future innovators.

Of course, because the relevant legal principles are still in flux, these guidelines represent merely one, general analysis of the legal landscape—please consult with an attorney regarding your precise plans.

1) Your two options: total control or total anarchy.

In the wake of the Napster decision, it appears that copyright law has foisted a binary choice on P2P developers: either build a system that allows for thorough monitoring and control over user activities, or build one that makes such monitoring and control completely impossible. This conclusion stems from the Ninth Circuit's analysis of contributory and vicarious liability in the Napster case.

As discussed above, contributory infringement requires that you have "knowledge" of, and "materially contribute" to, someone else's infringing activity. In most cases, it will be difficult to avoid "material contribution"—after all, if your system adds any value to the user experience, you probably have "materially contributed" to any infringing user activities. So the chief battleground on contributory infringement will likely be on the "knowledge" issue. The Napster court's analysis suggests that once you receive notice that your system is being used for infringing activity (e.g., a cease and desist letter from a copyright owner), you have a duty to "do something" to stop it. What might that "something" be? Well, it will be limited by the architecture of your system, but may ultimately be decided by a court. So, in order to avoid the unpleasant surprise of a court telling you to re-engineer your technology to stop your infringing users (as happened to Napster), you can either include mechanisms that enable monitoring and control of user activities (and use them to stop allegedly infringing activity when you receive complaints), or choose an architecture that will convince a judge that such monitoring and control is impossible.

The Napter court's vicarious liability analysis also counsels for either a total control or total anarchy approach. Vicarious liability requires that you "control," and receive "benefit" from, someone else's infringing activity. The "benefit" element will be difficult to resist in many P2P cases—so long as the system permits or enables the sharing of infringing materials, this will serve as a "draw" for users, which can be enough "benefit" to result in liability. So the fight will likely center on the "control" element. The Napster court found that the right to block a user's access to the service was enough to constitute "control." The court also found that Napster had a duty to monitor the activities of its users "to the fullest extent" possible. Accordingly, in order to avoid vicarious liability, a P2P developer would be wise to either incorporate mechanisms that make it easy to monitor and block infringing users, or choose an architecture that will convince a judge that monitoring and blocking is impossible.

2) Better to sell stand-alone software products than on-going services.

As discussed above, vicarious liability is perhaps the most serious risk facing P2P developers. Having the power to terminate or block users constitutes enough "control" to justify imposing vicarious liability. Add "financial benefit" in the form of a business model that depends on a large user base, and you're well on your way to joining Napster as a vicarious infringer. This is true even if you are completely unaware of what your users are up to—the pairing of control and financial benefit are enough. Of course, most "service" business models fit this "control" and "benefit" paradigm. What this means is that, after the Napster decision, if you offer a "service," you may have to monitor your users if you want to escape liability. If you want to avoid monitoring obligations, you'll have to give up on "control." It's time to set aside all the lessons you've learned about the importance of "relationships" in the New Economy. If your system may be used for infringement,

and this capability is a "draw" for users, you've already crossed the "benefit" threshold. In order to avoid vicarious liability for those infringing uses, you will need to give up any "control" over users.

Vendors of stand-alone software products may be in a better position to resist monitoring obligations and vicarious infringement liability. After Sony sells a VCR, it has no control over what the end-user does with it. Neither do the makers of photocopiers, optical scanners, or audio cassette recorders. Having built a device with many uses, only some of which may infringe copyrights, the typical electronics manufacturer has no way to "terminate" end-users or "block" their ability to use the device (but look out for those shrinkwrap software license terms permitting unilateral vendor termination). Not coincidentally, these manufacturers also typically don't get sued (at least not yet) by copyright owners. The key here is to let go of any control you may have over your users—no remote kill switch, contractual termination rights, or other similar mechanisms.

3) Can you plausibly deny knowing what your end-users are up to?

Assuming that you have escaped vicarious infringement by eliminating "control" or "financial benefit," there is still the danger of contributory infringement. To avoid liability here, you will need to address whether you knew, or should have known, of the infringing activity of your users. Have you built a level of "plausible deniability" into your product architecture and business model? If you promote, endorse, or facilitate the use of your product for infringing activity, you're asking for trouble. Similarly, software that sends back usage reports may lead to more knowledge than you want. Instead, talk up all the great legitimate capabilities, sell it (or give it away), and then leave the users alone. Again, the choices are total control, or total anarchy (see #1 above).

There are other good reasons for designing deniability into your product or system. First, it protects your users and, depending on your architecture, hosts or nodes as well. If you're not collecting information about what they're doing, no one can get that information from you. That's important for reasons that have little to do with copyright infringement. By not collecting user information, peer-to-peer networks can promote free speech and privacy. Remember the FBI's "Library Awareness Program"? Don't make yourself a target for subpoenas if you don't have to.

4) What are your substantial noninfringing uses?

If your product is intended to work solely as a mechanism for copyright piracy, you're asking for legal trouble. More importantly, you're thinking too small. Almost all peer-to-peer systems can be used for many different purposes, some of which the creators themselves fail to appreciate. So create a platform that lends itself to many uses, or, to paraphrase William Gibson, let the street find its own uses for things. For example, if you're developing a file-sharing system or distributed search

engine, support all file types, not just MP3 or Divx files. Actively, sincerely, and enthusiastically promote the noninfringing uses of your product. And don't promote any infringing uses.

The existence of real, substantial noninfringing uses will increase the chances that you can invoke the "Betamax defense" if challenged in court. As discussed above, however, it is worth noting that this defense will only help you until the copyright owner delivers a "cease and desist" letter notifying you of specific infringing activity. At that point, the "Betamax defense" may evaporate, and may leave you with an obligation to make a reasonable effort to stop the infringement. What this means will depend on the architecture of your system and the whims of the court.

5) Disaggregate functions.

Separate different functions and concentrate your efforts on a discrete area. In order to be successful, peer-to-peer networks will require products to address numerous functional needs—search (e.g., OpenCOLA), security (e.g., Intel's security toolkit), dynamic file redistribution (e.g., Freenet), to take a few examples. There's no reason why one entity should try to do all of these things. In fact, the creation of an open set of protocols, combined with a competitive mix of interoperable, but distinct, applications is probably a good idea, from a product-engineering point of view.

This approach may also have legal advantages. If Sony had not only manufactured VCRs, but also sold all the blank video tape, distributed all the TV Guides, and sponsored clubs and swap meets for VCR users, the Betamax case might have turned out differently. Part of Napster's downfall was its combination of indexing, searching, and file sharing in a single piece of software. If each activity is handled by a different product and vendor, on the other hand, each entity may have a better legal defense to a charge of infringement.

A disaggregated model, moreover, may limit what a court can order you to do to stop infringing activity by your users. For example, if a search engine that trawls the P2P network space were to be held contributorily or vicariously liable for facilitating access to copyrighted works, the search engine company would not be able make any changes to an anonymized, secure file-sharing product that was developed by a different company. As the Napster court recognized, you can only be ordered to police your own "premises"—the smaller it is, the less you can be required to do.

Finally, certain functions may be entitled to special protections under the "safe harbor" provisions of the Digital Millennium Copyright Act ("DMCA"). Search engines, for example, enjoy special DMCA protections. Thus, the combination of a P2P file sharing application with a third party search engine might be easier to defend in court than Napster's integrated solution.

6) Don't make your money from the infringing activities of your users.

Avoid business models that rely on revenue streams that can be directly traced to infringing activities. For example, if you are developing a peer-to-peer auction system, do not take a percentage cut on transactions completed through the system. To take another example, a peer-to-peer file sharing system that includes a payment mechanism might pose similar problems, if the system vendor takes a percentage cut of all payments, including payments generated from sales of bootleg Divx movie files.

Of course, in the wake of the Napster decision, the mere fact that infringing material may act as a "draw," thus increasing your user base, might be enough to trigger vicarious liability. Nevertheless, there is nothing to be gained by building your business on a "financial benefit" even more directly linked to infringing activity by users—you'll only be making it that much easier for copyright owners to shut you down.

7) Be open source.

In addition to the usual litany of arguments favoring the open-source model, the open source approach may offer special advantages in the peer-to-peer realm. It may be more difficult for a copyright owner to demonstrate "control" or "financial benefit" with respect to an open source product. After all, anyone can download and compile open source code, and no one has the ability to "terminate" or "block access" or otherwise control the use of the resulting applications. "Financial benefit" may also be a problematic concept where the developers do not directly realize any financial gains from the code (as noted above, however, the Napster court has embraced a very broad notion of "financial benefit," so this may not be enough to save you). Finally, by making the most legally dangerous elements of the P2P network open source (or relying on the open source projects of others), you can build your business out of more legally defensible ancillary services (such as search services, bandwidth enhancement, file storage, file meta-data services, etc.).

8) Do not be a direct infringer: make and store no copies.

This one may be obvious, but remember that if you make or distribute any copies (even if only in RAM) of copyrighted works, you may be held liable as a direct infringer. In that case, many of the defenses discussed here will not be available to you. The court will not be interested in "control" or "knowledge" or "financial benefit" or "material contribution." If you made copies, you're probably liable for infringement. Of course, this shouldn't be a problem for most P2P developers, since the great insight of peer-to-peer architectures is that the actual resources being shared need not pass through any central server. Nevertheless, be careful where caching or similar activities are concerned.

9) Do not build any "circumvention devices" into your product.

Avoid incorporating into your product any technology designed to circumvent a protection measure meant to protect the rights of copyright owners. For example, do not incorporate any code into your product that is intended to crack the CSS encryption system used on DVDs, bypass the protection scheme used on Sony Playstation video games, or strip the "no copy" flags out of streaming RealAudio files. Whatever you may think about the merits of this work, leave it to others. You'll have enough worries without adding circumvention liability to the list.

10) Don't use someone else's trademark in your name.

This tip does not relate to copyright, but rather trademark law. It's also good common sense. Many early peer-to-peer products are attempting to capitalize on the notoriety of Napster by incorporating portions of the Napster name into their products—Napigator, OpenNap, and AIMster, to name a few. Napster has shown itself to be a strong defender of its trademark rights, having threatened legal action against a variety of unauthorized Napster merchandise vendors. Gnutella may also be a dangerous name to appropriate, as it is not clear who owns it (AOL?), and whether the owners of the Nutella trademark may at some point assert their own trademark rights. And remember that AppleSoup, for example, was forced to changed its name to FlyCode after receiving pressure from Apple Computer.

Good luck, and change the world!

* * *

About the Author: Fred von Lohmann is a senior staff attorney with the Electronic Frontier Foundation, specializing in intellectual property issues. In that role, he has represented programmers, technology innovators, and individuals in litigation against every major record label, movie studio, and television network (as well as several cable TV networks and music publishers) in the United States. In additon to litigation, he is involved in EFF's efforts to educate policy-makers regarding the proper balance between intellectual property protection and the public interest in fair use, free expression, and innovation.

Footnote

* Acronym for "I am a lawyer," to distinguish from the common "IANAL" ("I am not a lawyer") that appears on Slashdot and other online forums.