

Distributed File Sharing: Barbarians at the Gates?

Michael Macedonia, US Army Stricom

The Web is now a vast storehouse of musical treasures—a startling discovery that has swept through many households and college campuses this year thanks to the emergence of a single distributed file sharing (DFS) program: Napster. Traditionally, the entertainment industry makes the bulk of its earnings by controlling the rights and distribution of its performers' intellectual property. DFS provides a major challenge to the entertainment business's rules because it lets consumers link to locally stored music files from around the world, disseminating this content in digital form for next to nothing. Ironically, artists such as Metallica and Dr. Dre, whose music celebrates social anarchy, must now worry about the emergence of this *digital* anarchy.

Not that they're alone. Napster's enormous legal implications are being vigorously dissected in the press, the legislature, and courtrooms at every level. Although law and business have acquired a decade's experience in dealing with the Web's effects on their domains, it will take them many years to sort out the consequences of a maturing DFS technology. In a recent issue of *Fortune*, Amy Kovar quotes Intel Chairman Andy Grove as observing that "The whole Internet could be re-architected by Napster-like technology" ("Napster: The Hot Idea of the Year," <http://library.northernlight.com/>

LH20000616010000693.html?cb=13&sc=0#doc_).

TECHNOLOGIES CONVERGE

Thanks to Napster, DFS has caused quite a stir by escaping the boundaries of the office and campus to go global via the

Internet. Yet file-sharing software is not new. Sun, for example, developed its Network File System (NFS) in the 1980s to let users store and share files among LAN servers and workstations. NFS lets files on remote systems appear as if they reside in a local system's directory. For wide area use, however, most files must usually be retrieved from a large server, such as a mainframe, through a relatively slow Internet connection using file transfer protocol (FTP).

An Internet for the masses

Several factors converged to change this model and make a program like Napster possible. First, the arcane Internet morphed into the user-friendly World Wide

Web. Melded to graphical browser clients the nontechnical could comprehend, the Internet ceased being the exclusive domain of scholars, students, and the military.

Second, the increasing movement toward broadband communications in the home, particularly with the emergence of digital subscriber line (DSL) and cable modems, gave home consumers high-speed Internet connections that rivaled those previously restricted to major corporations. With DSL, for example, a 1-Mbyte file that would take two minutes to download on a conventional modem can be acquired in about eight seconds.

Digital compression

Third, digital media and the software to play it developed with explosive speed. Compact disks led the charge, digitizing music with Eight to Fourteen Modulation (EFM). This system encodes data when it's recorded and decodes it when it's played. During encoding, EFM translates a series of 8 data bits into an arbitrary set

Will the MP3 format and Napster's distributed file sharing bankrupt music publishers—or help them reach even more consumers?

of 14 data bits, which are returned to the original 8-bit format during decoding. This modulation serves the two-fold purpose of implementing a rudimentary type of error correction while increasing disc storage efficiency by reducing the number of pits.

In addition to CDs, proprietary formats such as Microsoft's wave audio files also became popular, but had limited utility because they produced files so large it took 32 Mbytes to store a three-minute song. Thus, the digital migration stalled until new compression technologies shifted it into overdrive—specifically the lossy compression technique called Motion Picture Equipment Group (MPEG) Audio Layer 3, better known as MP3 (<http://www.iis>

fhg.de/amm/techinf/layer3/index.html). Developed by the Fraunhofer Institute for Integrated Circuits and the University of Erlangen, MP3 can compress CD-quality audio by a factor of 10, reducing that three-minute song file to 3 MBytes.

Today, tools known as rippers—easily available as shareware or freeware—can convert a music CD's audio tracks to MP3 format, essentially "ripping" them off the disc. Digital audio players include well-known software like WinAmp (<http://www.winamp.com/>) and any number of proprietary packages bundled with PC sound cards.

More speed, more space

Fourth, the breakneck pace of personal-computer development has given us consumer-grade PCs with fast processors and huge storage capabilities. Generally, faster computers mean faster network processing and the ability to run computationally expensive, multimedia content such as MP3s and digital video. Most new PC processors operate at 500 MHz or more—a speed that's likely to increase rapidly given that the heated competition between chip makers AMD and Intel has already broken the 1-GHz barrier.

PC data storage capacity has shown similar progress, with 60-Gbyte hard drives now available for less than \$300. Less than a decade ago, nothing short of a mainframe could store this much data. Portable media are equally cheap, with recordable CD drives now costing what floppy drives once did, and the blank CDs needed to feed them priced at a mere \$.50. Yet each disk holds 640 MBytes of data—or more than 200 MP3 files.

INNOVATION AND DIFFUSION

All these technologies were thriving in the marketplace last year. It took a 19-year-old computer science student at Boston's Northeastern University to synthesize them into a revolutionary new mix.

A simple plan

Shawn Fanning just wanted an easy way to find and share MP3 files. Nicknamed "Napster" in high school, Fanning conceived a program to store

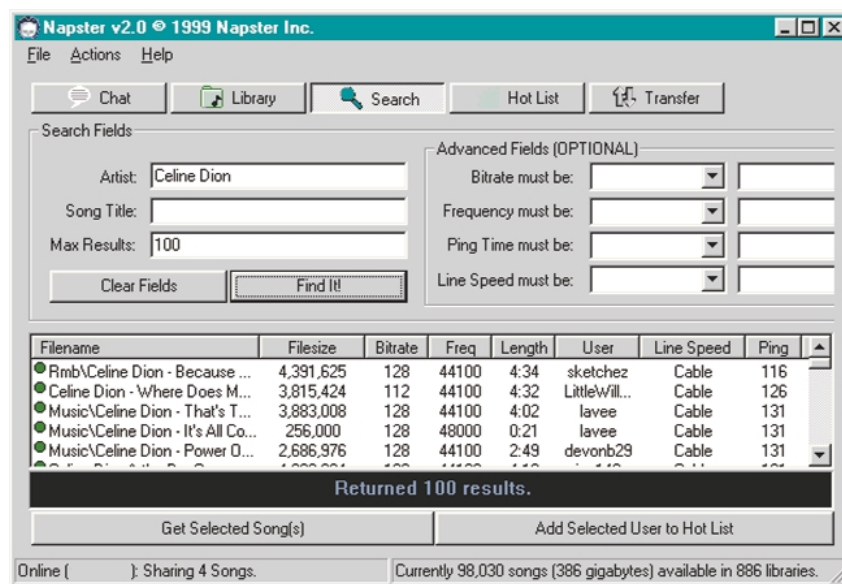


Figure 1. Napster makes it easy to find and download music files stored on hard disks around the world.

and access MP3 files. His creation takes advantage of the extra space available on most home computers' hard drives—and their constant or near-constant connection to the Internet.

Fanning turned his idea into a company, Napster Inc. (<http://www.napster.com/>). Napster's very simple architecture works by reserving a portion of each user's local hard drive for storing MP3 files ripped from audio CDs. The client software then advertises the files to a group of Napster servers. Users browse a huge directory on the servers and, when someone wants to retrieve a particular song, the server points that person to any of several user hard drives that contain the desired file. To help estimate download times, Napster also tells the user what kind of network connection the file's owner has.

Developing a convenient way to share files has proven Fanning's greatest contribution. Napster, like successful services such as America Online, is not the most sophisticated creation, but is easiest to use and most popular. AOL President Bob Pittman seemed to acknowledge this parallel between Napster and his company when he observed at the Electronic Entertainment Expo in Los Angeles that "The recent legal hoopla over MP3s and

digital music doesn't show that consumers want to be thieves, it shows that they want to listen to music on their PCs" (<http://europe.cnn.com/2000/TECH/computing/05/12/aol.keynote.idg/>).

Indeed they do. According to Kovar, Napster has attracted 10 million users in a mere 10 months.

Napster's dark side

Napster has its pathologies, however. Indiana University found that Napster downloads of MP3 files accounted for 60 percent of the school's network traffic (<http://www.fortune.com/fortune/2000/03/20/nap.html>).

Napster use also raises security questions since the user must share part of the local disk while online. Worse, others have already crafted an application, "Wrapster" (<http://members.fortunecity.com/wrapster>) that archives any kind of files—such as a Perl script or even a virus—inside a legitimate MP3 file.

Moreover, since the Napster servers do not actually store MP3 files, the company argues it is not violating copyright laws. Beyond the legal gray area of whether the company is facilitating copyright infringement, there is no evidence of a business model at work because Napster charges nothing for its service.

Gnutella and other Napster variations

A more anarchic version of Napster's file sharing model, Gnutella (<http://www.gnutella.wego.com>), seeks to eliminate the middleman altogether. Developed by an employee of AOL subsidiary Nullsoft (<http://www.nullsoft.com/>), Gnutella was accidentally posted to the Web and withdrawn days later. Yet unauthorized copies continued to spread while enthusiasts—with the original author's help—reverse engineered the Gnutella protocol in Linux fashion and created several new clients, including the highly popular Gnut (<http://www.umd.edu/~jjp/>).

Gnutella eliminates the central server for files and file directories. Instead, clients propagate search requests among themselves. When you retrieve the client software, you initiate your searches first through well-known sites published on mail lists or the Web. This method, similar to the spanning-tree algorithm used for

network bridging, uses time to live (TTL) counters to prevent loops. It also allows your client to quickly identify thousands of other clients and locate a file, whether MP3, video, or some other format (http://members.tripod.com/~retep/articals/file_sharing_programs.html).

Napster, like successful services such as America Online, is not the most sophisticated creation, but is easiest to use and most popular.

Gnutella pays a penalty for its independence from a central server, however. It's not as easy to use as Napster. Worse, its searches consume substantial bandwidth since the request messages replicate rapidly and widely over the Web, much like a message with a rogue visual basic script. Moreover, Gnutella lacks a business model because there is no business. Even the clients are free.

Another Napster alternative, the experimental DFS system Freenet (<http://freenet.sourceforge.net>), is also decentralized. Freenet attempts to reduce bandwidth requirements by moving files among hosts to implicitly balance traffic, similar to Web caching. In addition, Freenet encrypts the files on each user's hard drive so that even the user does not know which files are locally present or shared. In this case, ignorance is theoretical bliss with regard to copyright law.

OUT OF CHAOS, MORE CHAOS

Until DFS grows a business model, it won't make anybody rich. So, on the one hand, the entertainment companies may have little to worry about.

On the other hand, we live in Internet time, so a succession of DFS clients and related innovations will likely continue to outpace the law. Moreover, DFS is immune to political boundaries—much like Internet gambling, which has moved to off-shore domains. These factors imply that life will only get more complicated for those whose livelihood depends on pumping CDs and DVDs into stores.

The entertainment industry may try to counter the parasitic effects of DFS with digital watermarking, copy protection schemes, and encryption. However, these methods are peripheral to the issue of digital product distribution. With the exception of AOL/Time Warner, leading entertainment companies such as Universal Music, Sony, and Bertelsmann do not have a major stake in the Internet. Neither do they control the personal computer market—a major concern considering that in five years the average PC will likely have the storage capacity to archive a music publisher's entire library.

However, the entertainment industry does appear to be realizing, finally, that the Web is swarming with young and affluent consumers. An obvious strategy, given the industry's lack of leverage in the PC sector, is to move consumers away from personal computers.

The Trojan horse that decides this campaign may well be the inexpensive game console. Sony's Playstation II is a genuine multimedia machine complete with AV ports, DVD, a 3D sound system—and a broadband connection. Microsoft's X-Box boasts similar features and ups the ante with a hard disk ideal for storing MP3 files. Not only are these boxes proprietary, their network portals will likely be as well.

Given technology's relentless pace, a Palm X or Nokia 3G phone may become the ultimate MP3 platform of choice. I'm putting my money on a closed, Web-based system, however. Regardless, with AOL/Time Warner—the world's largest media-Internet-cable-browser-e-mail company—in this race, we should be guaranteed an exciting ride. Let the games begin. *

Michael Macedonia is chief scientist and technical director of the US Army Simulation, Training, and Instrumentation Command, Orlando, Fla., (<http://www.stricom.army.mil/stricom/aboutus/macedonia.html>).

Michael Macedonia, Stricom, 12350 Research Parkway, Orlando, FL, 32826-3276, macedonia@computer.org