

puters, viewing those changes in technological and historical perspective. I teach courses entitled “Law in the Internet Society,” “Computers, Privacy, and the Constitution,” and “Perspectives in Modern Legal Thought.”

7. I have published widely in these fields, and in the field of American legal history, specializing in the development of law in English-speaking North America from the inception of colonization through the American Revolution. A copy of my *curriculum vitae*, including a list of my publications, is attached hereto as Exhibit A.

8. Since 1994 I have served *pro bono publico* as General Counsel of the Free Software Foundation. I have been a member of the Foundation’s Board of Directors since 1999.

The Free Software Foundation

9. The Free Software Foundation (“FSF”) was founded in 1985 by Richard M. Stallman, who remains its President. The Foundation is a 501(c)(3) organization incorporated under the law of the Commonwealth of Massachusetts, with its primary place of business in Boston, Mass.

10. The FSF’s mission is to encourage the creation and distribution of computer programs, technical documentation, and other related materials that can be freely copied, modified and redistributed by their users. FSF refers to such computer programs as “free software,” where the word “free” refers to freedom, not to price. FSF believes that by giving all users the right to copy, modify, and redistribute software, the ethical obligation to maintain freedom of thought is honored. Software that can be freely shared is also of inherently higher quality, because everyone who uses the software can experiment with improvements, and can fix mistakes that are discovered. Because anyone who makes such fixes is allowed to distribute those fixes or improvements, the quality of the software increases exponentially over time, and all programs can be reused for new purposes.

11. FSF supports the development and distribution of free software in two basic ways: it writes and distributes free software of its own, and it helps others to write and distribute such software, primarily through the development and publication of the copyright licenses and associated legal materials that facilitate this hitherto-unusual means of software production.

12. Before the creation of the Foundation, Mr Stallman had begun designing and implementing an entire free software operating environment usable on all computers from the most basic facilities to the most advanced applications. He called that planned system “GNU.” Over the course of the 1980s many components of GNU were written and tested, some by Mr Stallman himself and many others by programmers who contributed to the projects and assigned their copyrights in their portions of the programs to the Free Software Foundation.

13. In the early 1990s a young Finnish programmer named Linus Torvalds began to assemble the innermost portion of a computer operating system, known in the trade as the “kernel,” first as a personal learning exercise and then as a cooperative project over the Internet, ultimately involving thousands of other volunteers. Mr Torvalds called his kernel “Linux,” and he designed it to work compatibly with the other parts of the GNU system designed and incrementally implemented by Mr Stallman and FSF. The result was a system that combined Mr Torvalds’ Linux kernel with the other GNU components to make a system that FSF calls GNU/Linux, and which is widely but misleadingly called “Linux” in general parlance.

14. Throughout the 1990s the GNU/Linux operating system became explosively popular with technically-sophisticated users and businesses around the world. The production model of free modification and redistribution permitted rapid development of thousands of applications for the system, and its compatibility with the Unix operating system originally designed and implemented by AT&T made possible the ready adaptation or “porting” of most applications designed for Unix, itself very widely used in technical, academic, and scientific environments. It has been widely reported that GNU/Linux is now the fastest growing operating system in the world for “server computers,” which are those computers that perform the tasks required by large networks: file sharing, World Wide Web publication, etc. GNU/Linux is now also used in the smallest computers in the world, including “personal digital assistants” or “palmtop” computers. Because, in compliance with the terms and conditions of the GPL, anyone may freely copy any or all of the programs contained in the system, centralized sales and use figures do not exist, but even a conservative estimate of the number of computers using the operating system and associated application programs throughout the world would reach

into many tens of millions.

15. Although high technical quality and reliability is certainly partly responsible for the success of GNU/Linux, the legal institutions that facilitate this apparently counter-intuitive phenomenon of large-scale non-hierarchical production successfully competing against global corporations are of even greater importance. These institutions depend on a simple but far-reaching employment of copyright law through the GNU General Public License, under which much free software is distributed.

The GNU General Public License

16. The GNU General Public License (“GPL”) is the legal heart of the free software movement. The goal of the GPL is to use copyright law to create a “commons,” a collection of shared resources to which anyone can add, and from which anyone can borrow freely, but from which nothing can be permanently removed. This concept, of using copyright to create a commons, rather than a domain of exclusive ownership, is sometimes called “copyleft,” and the GPL is an example of one form of “copyleft license.” A copy of the current version 2 of the GNU GPL, first published in 1991, is attached hereto as Exhibit B.

17. Free software is not in general in the public domain. If contributors to free software projects were placing their code in the public domain, it could be immediately incorporated by others into proprietary, non-free projects, from which those appropriators could derive value without returning anything to the commons. Instead, free software is copyrighted, but through the terms of the GPL, copyright is used to protect the common interest without excluding anyone from the rights to execute, copy, modify, and redistribute that are of the essence to free software.

18. The GPL is a very simple form of copyright license, as compared to other current standards in the software industry, because it involves no contractual obligations. Most software licenses begin with the exclusive rights conveyed to authors under copyright law, and then allow others access to the copyrighted work only under additional contractual conditions. The GPL, on the other hand, actually *subtracts* from the author’s usual exclusive rights under copyright law, through the granting of unilateral permissions. When a work of copyrighted software is released under the GPL, all persons everywhere observing its terms are unilaterally

permitted all rights to use, copy, and modify the software. Because these permissions are unilaterally given, users who wish only to use the software themselves, making copies for their own use, or who wish only to make derivative works for their own use, do not have to “accept” the license, because they have no reciprocal obligations under it.

19. If a user wishes to redistribute software she has received under the GPL, whether in modified or unmodified form, the license permits that activity as well. Here, however, the permission is qualified by three primary conditions:

- Redistribution must itself occur under GPL and only GPL, with no additional license conditions. (See Exhibit B, §2(b));
- Redistribution must include “source code,” the human-readable form of computer programs that allows programmers to understand and modify computer programs for themselves, as opposed to “object code,” which is the “machine language” version of computer programs that is very difficult for programmers to understand or modify. (See Exhibit B, §3(a)); and
- Redistribution must include a copy of the GPL, so that users are aware of their rights to use, copy, modify and distribute, and so that anyone engaged in redistribution is also aware of the conditions under which redistribution is permitted. (See Exhibit B, §1).

20. As a result of these conditions on redistribution, the GPL achieves the goal of creating a commons. Anyone can copy and modify program code released under the GPL, but no one can combine that program code with any other code and then release the combination on non-GPL terms. Anyone who contributes program code to a GPL-based programming project knows that her contribution will remain freely available for others to use, fix and improve, but that no one will be able to exclude others from having the same rights. The GPL uses copyright doctrine to achieve the result of the principle that we should all “share and share alike.” Thus anyone who combines GPL-licensed software with other program code must release the combined work under GPL, and must provide the source code for the entire derivative work.

21. Because anyone in possession of a program released under GPL must be in actual possession of the license itself, the licensor is entitled to presume that anyone engaged in redistribution is actually on notice of the only terms on which redistribution is permitted. Redistribution on any other terms is intentional violation of the GPL. (See Exhibit B, §5).

22. The GPL is specifically designed to be a license for decentralized distribution, in which everyone can share programs and improvements with anyone else. This means that program code can cross national borders and otherwise propagate in uncontrolled ways. For this reason, the GPL makes special provision for dealing with the consequences of license violation. Under §4, any licensee who violates the GPL loses his right of distribution, until such time as that right is restored by affirmative act of the copyright holder. The distributees of that licensor, however, retain their rights under the license, including their rights of distribution. (See Exhibit B, §4).

23. The FSF is by no means the only licensor of programs under the GPL. FSF accepts copyright assignments of some programs for release under GPL, as part of its mission to facilitate the employment and diffusion of free software. The IBM Corporation, for example, not only releases some programs under the GPL, but also assigns copyright to FSF in some of the programs it so releases, for the purpose of empowering FSF to enforce the GPL against license violators. But many other authors of programs choose to release their works under GPL while retaining the ownership of their copyrights: the Linux kernel itself, for example, is owned by its authors; FSF has no significant ownership of copyrights in the primary versions of the Linux kernel. FSF's own enforcement practices in dealing with GPL violations, however, are widely followed in the community of free software and (as it is sometimes called) "open source," a phrase which refers to the requirement that source code be provided to facilitate understanding, modification, and redistribution by all users. FSF's enforcement practices are relevant to the community at large because FSF is the author of the GPL, because it has a large inventory of free software which it manages on behalf of the community as a whole, and because its strong engagement with the ethical as well as commercial value of free software has given it a special position of trust within the community.

24. In my role as General Counsel of the Foundation, I have been primarily responsible for all world-wide enforcement activity in defense of the GPL since 1994. I have been involved in all significant cases in which FSF itself enforced the license with respect to software whose copyrights it held, and I have provided extensive advice to authors of other software released under the GPL with respect to the enforcement of their rights and protection of the integrity of the license. In those activities, I have found §4 of the GPL to be absolutely essential to the conduct of our enforcement strategy.

25. FSF's policy with respect to GPL violations is to secure compliance, not damages. When a party has violated GPL, and the violation is called to our attention (which happens on the average some dozens of times each year), we inform the party in violation of its responsibilities, and advise it on the steps necessary to come into compliance. It is our practice that once a party has taken steps to comply, and has entered into confidence-building measures to ensure that future non-compliance will be avoided wherever possible, and rapidly discovered and remedied where inadvertently reproduced, distribution rights under GPL §4 are restored on a cooperative non-judicial basis. In this fashion, I have secured compliance with the license in dozens of cases over the past decade, and have never had to resort to judicial measures of mandatory enforcement. Without the leverage provided by §4, however, parties would resort to repetitive partial compliance, "capable of repetition but evading review," in language the Supreme Court has applied to a different sort of situation, substantially if not overwhelmingly complicating the task of securing reliable compliance with the license.

MySQL

26. I understand that MySQL AB was formed by the primary authors and is a copyright holder of the program MySQL, which is a particularly important component among the many thousands of components of the free software system. MySQL is a "database engine," which means that it organizes and manages access to large quantities of "tabular" data.

27. Database tables would include the transactions in a store, the reservations in an airline reservation system, the addresses in a mailing list, or the personnel records of the employees of a company. Each "record"

in such a database can be thought of as the “row” of a table, in which a single sales transaction, for example, contains several “fields,” or “columns” (the item number of the item sold, the quantity, the price, the shipping charge and tax applied, and so on). The “SQL” in the name stands for “structured query language,” which is the historically-conventional phrase used to describe database engines that allow rows in the table to be searched for by a complex set of restrictions on the fields (“Give me all the transactions in which someone in zipcode 10027 bought five or more of item #3116 in color green or blue during the months of April and May 2000,” for example).

28. The “engine” which manages the creation and searching of such tables, particularly when they involve hundreds of thousands or millions of entries, is a critical building block of larger software applications. Most of the “e-commerce” occurring on the World Wide Web, whether it involves making a purchase in a store, reserving transportation or accommodation, or tracking the movement of packages by freight companies, involves interacting with a database manager.

29. MySQL is the most popular and widely-used of free software database engines. Along with the program called Apache that manages websites, and the Perl and PHP “scripting languages” that allow web pages to accept and process user input, MySQL is part of the “platform” on which many thousands of individual applications have been constructed by businesses and other organizations, large and small, around the world.

30. MySQL AB engages in “dual licensing.” This means that it licenses a version of MySQL to be freely used, copied, modified and distributed by everyone under the GPL, and also makes versions of its program that are distributed to particular customers without the right of free distribution. Those who receive MySQL under the GPL, however, are not entitled to engage in “dual licensing.” Having received their copy of the program under GPL, they may freely modify and redistribute, but that redistribution, under GPL §2(b), must occur under the terms of the GPL, without any additional limitation. In particular, anyone who modifies MySQL must release that modified version in compliance with the GPL’s requirement that everyone who receives the program must receive “source code,” or all the materials required so that they themselves can

understand, share and improve the program in its modified form.

31. Progress Software Corp. is the distributor of program code called “Gemini.” Gemini is a “storage module” for a database engine. It performs the task of writing data into a new or existing database, which is a subcomponent of the tasks performed by the database engine as a whole. Gemini provides numerous features for any database engine into which it is inserted. Gemini can provide “crash recovery”: if a large database stored by Gemini is in the process of being updated when the computer system managing the data is suddenly halted, by a power outage or similar unexpected intervention, Gemini can assure that the data is uncorrupted. Gemini also provides “reversability,” in the sense that a change made in the database is not final when it is first made, and thus, if it turns out that the change is erroneous, until the table stored by Gemini is “committed,” the original pre-change data can still be recovered. But Gemini tables are also larger than tables stored by other storage managers, and the task of storing information in such tables is slower. For this reason, a database engine such as MySQL relies upon multiple storage modules, rather than only one. When a particular data generator decides how to construct a database for a particular task, it will instruct the database engine to choose among storage modules in order to achieve the optimal balance of size, speed, stability, reversability, and other properties. In any particular computer system that provides database applications, the MySQL engine will be a single program, usually called “mysqld.” The mysqld program on that system will have one or more storage modules “compiled” or “linked” into it as subcomponents, like chapters in a single book, and those particular chapters will have been included that serve the needs of the particular application or applications that will use database services.

The Current Dispute

32. In connection with the instant litigation, I have reviewed two versions of “NuSphere MySQL Advantage” distributed by Progress Software Corporation. I have conducted that review precisely as I would have conducted such an investigation in my role as FSF’s General Counsel in the event of a complaint of GPL violation. I personally conducted all activities hereinafter described, depending on my own personal knowledge of the practices of software manufacture and distribution, as well as my legal understanding of the

requirements of FSF's GNU General Public License.

33. NuSphere MySQL Advantage is a collection of programs, like a bookshelf containing many books, all of which have been adapted to work smoothly together. Among the separate works anthologized in NuSphere MySQL Advantage are the Apache web server program, the Perl and PHP scripting languages, and MySQL. A user who buys NuSphere MySQL Advantage is thus getting "one-stop" shopping: each of the programs involved is free software, released under a variety of free software licenses, and could each be gotten from multiple places in the network, usually at nominal cost. By paying a substantial price for NuSphere MySQL Advantage, however, the user gets everything all on one CD ROM, arranged to work smoothly together and to be installed using a simple interface that can be controlled from an ordinary web browser, and which will install the same way on machines running Microsoft Windows, GNU/Linux, or other operating systems for non-PC computers.

34. The first version of NuSphere MySQL Advantage I reviewed was labeled Version 2.2. The CD ROM containing the software associated dates with each file in the customary fashion, and by reviewing the dates I concluded that the CD ROM was "burned," or fixed in a tangible form, on June 21, 2001. I installed this copy of NuSphere MySQL Advantage on a computer running GNU/Linux, following the standard directions provided. The standard installation program created a `mysqld` program on my computer's hard drive, and once that program had been created I used a standard information request to inquire of that copy of `mysqld` whether it included Gemini storage capacity. The program told me that Gemini storage capacity was included in `mysqld`, and gave me certain technical particulars, for example the largest Gemini table that it could store.

35. I then checked whether the source code of the Gemini component of `mysqld` was available. I saw that it was not. I used standard tools to "recompile" `mysqld`, which means to rebuild the program from the parts provided, and discovered that the rebuilt `mysqld` did not include Gemini capacity. In a manual distributed in the portion of the CD only related to Windows users I found the statement that "'GEMINI' tables will be included in some future MySQL 3.23.X source distribution."

36. On the basis of this examination, based upon my expertise and prior experience as a computer

programmer, I concluded that the version of NuSphere MySQL Advantage under review violated the requirements of the GNU GPL. The mysqld program as distributed was covered by the GPL. It had been combined with program code from Gemini to provide those storage types, and had been “compiled” into the “machine language” copy of mysqld distributed to buyers. But the source code for the portion of mysqld that performed Gemini storage had been withheld, which violated GPL §3. This omission was not inadvertent, as was shown by the comment in the Windows version of the MySQL documentation, which said only that source code would be provided at a later date. Given my knowledge of the practices of programming and the requirements of the GPL, I concluded that the license violation was intentional. Whether intentional or not, any violation of the GPL results under §4 in a termination of the right to redistribute.

37. I then reviewed a copy of Version 2.3.1 of NuSphere MySQL Advantage. The dates in the filesystem allowed me to conclude that the CD ROM was “burned” on October 1, 2001. Installing as ordinarily directed I again verified that the version of mysqld placed on my hard drive contained Gemini storage capacity. I reviewed the source code directories and found that the source code of the Gemini module was fully available. I then verified that I could “compile” or rebuild mysqld from the available parts so that it matched the version that had been distributed in “machine language” on the disk. The statement about postponed source code availability had been removed from the Windows manual.

38. Under GPL §4, I conclude, Progress Software Corp. lost the right to distribute MySQL when it distributed NuSphere MySQL Advantage in a fashion that violated GPL.

I declare under penalty of perjury and upon personal knowledge that the foregoing is true and correct.

Dated: February 26, 2002
New York, New York

Eben Moglen