

Closed source fights back

In May 2003, SCO, a vendor of the Linux operating system, sent a letter to its customers. Amongst other things, it stated:^[1]

We believe that Linux is, in material part, an unauthorized derivative of UNIX.

What would make them do that?

The action wasn't completely unexpected: in March, SCO had filed a suit against IBM for giving away trade secrets.^[2] In that complaint, they made a number of accusations against IBM, including the claim that they intended to use Linux to kill UNIX.

For most people in the industry, Linux *is* UNIX. Admittedly it was developed independently of UNIX, but the similarities are marked and intentional. So how can Linux kill UNIX? That depends on your interpretation of the word UNIX. There are a number of ways to view UNIX: you can take the genealogical approach ("It was derived from the same code base"), the standards-based approach ("It complies with UNIX 03"), the legal approach ("It has been certified as being UNIX 03"), or the technical approach ("It looks like UNIX, smells like UNIX, so it must be UNIX"). Based on the last criterion, many people would say that Linux is indeed UNIX.

UNIX genealogy

The UNIX family has hundreds of members, some of which have been disowned by their parents. Éric Lévénez maintains a UNIX family tree^[3] which, printed out in miniscule print, is about 10 feet long. In brief, though, there have been three main versions based on the original source code:

- *Research UNIX* was the original UNIX developed at Bell Labs.
- *BSD* (for *Berkeley Software Distribution*), also known as *Berkeley UNIX*, is the version developed at the University of California in Berkeley. It was the predominant UNIX in the 1980s, helped by the development of the Internet and its adoption by a number of workstation manufacturers, notably Sun Microsystems and Digital Equipment Corporation.
- In the early 80s, AT&T developed *System III* and then *System V*. They persuaded a number of workstation manufacturers to switch from Berkeley UNIX to System V.

During the 80s, a number of things happened:

- UNIX became less interesting as a research subject. The original Bell Labs team moved on to other projects. The Computer Sciences Research Group (CSRG) in Berkeley prepared to close down by the early 90s.
- Vendors using UNIX System V found it necessary to adapt their products more and more, both for performance reasons and for market differentiation. Only a few years after the release of System V.4, almost nobody was running stock System V.4. SCO was running a strongly modified version of System V.3, and the hardware manufacturers all had their own versions.
- Some vendors, led by Digital Equipment Corporation and IBM, became sufficiently dissatisfied with the direction in which System V was going that they started their own version, OSF/1. Initially OSF/1 was supposed to be free of AT&T code, but they never achieved that goal.

What was wrong with System V? To quote a button depicted on the front of Peter H. Salus' "A Quarter Century of UNIX", published in 1994: "System V does everything Unix does, only not as well". More technical reasons were a lack of good memory management, almost non-existent support for multiple processors, and an installation procedure that paid more attention to licensing dozens of individual components than it did to ease of use. In January 1992, the German magazine iX had planned a review of no less than five different versions of UNIX System V.4.^[4] At that time, all versions of UNIX sold as "System V.4" ran on the PC platform.^[5] Over several weeks they tried multiple machines, but they were able to get only one of the five systems to boot. They gave up on the review.

Hardware vendors had their own problems. The initial release of System V.4, the agglomeration of System V.3, Mi-

Microsoft's (later SCO's) XENIX and 4.3BSD, was so buggy that those vendors who based their UNIX offerings on System V.4 had to spend up to two years debugging before they could ship a product. They also adapted the code to their systems, fixing technical weaknesses in the process, addressed the administrative issues, and gave the product a distinct name. It's not surprising that the market share of stock System V began to dwindle. The most recent System V implementation was UnixWare, created by USL and Novell in 1992. It successfully addressed the weaknesses mentioned above, but it was too late. Stock System V had started to gradually disappear even before Linux and the free BSD versions became widespread.

Who owns UNIX?

We've already seen that "UNIX" means many things, some of which associated with a monetary value. In particular, the trade mark, the source code base and associated patents are considered valuable. Each of these belongs to a different organization:

- The Open Group^[6] owns the UNIX trade mark. The main use of the trade mark is for certification. Certified products are not necessarily UNIX in the genealogical sense. Theoretically Linux could be certified as UNIX. The most recent standard, the Single UNIX Specification Version 3, is also called UNIX 03.^[7] It was announced in January this year. Currently no products are certified to this standard. The previous standard is the Single UNIX Specification Version 2, or UNIX 98. Only products from Compaq,^[8] Fujitsu, IBM and Sun are certified as UNIX 98.^[9] The Open Group has recently issued a statement^[10] in which they say: "SCO has never owned "UNIX"."
- In a letter to SCO,^[11] Novell asserts that it owns the UNIX patents and copyrights. There has been some discussion about this matter, but as of going to press Novell has not retracted that statement.
- SCO owns the UNIX System V code base.

It's interesting to note that of these three organizations, only SCO shows any opposition to Open Source. In addition to UNIX certification, the Open Group also provide Linux certification in accordance with the Linux Standards Base, and Novell has openly stated its support for Linux.^[12]

Why is SCO attacking Open Source?

SCO is acting as if they are the most important player in the UNIX game. In fact, they're relatively insignificant. The last time System V was licensed to a third party was in 1994, to HP for upgrading HP-UX to release 10. HP-UX has been based on System III and System V since its initial release in 1982, so this wasn't exactly a new customer. The last time before that was to Sun in 1991, to create SunOS 5.^[13] Their own products do not comply with the latest UNIX specification. SCO's most up-to-date product, UnixWare, is certified as conforming to the Single UNIX Standard Version 1, or UNIX 95.^[14] SCO OpenServer, based on UNIX System V.3.2 (introduced in 1988), is certified as UNIX 93.^[15] By all accounts they are not selling well. I don't know where I could buy UnixWare, apart from directly from SCO. They announced UNIX System V.5 in 1997, but so far the only implementation is UnixWare 7.^[16] Like many other Linux vendors, they had not found a way to be profitable distributing and supporting Linux. Their most unique asset was the System V code base, the basis of nearly all current commercial UNIX distributions.

But what are those distributions doing? They're gradually going away. There are only really four major UNIX vendors left: HP, IBM, SGI and Sun. All of them are actively developing Linux. SCO sees its revenue streams going away. You can't get revenue from free software.

On the other hand, Linux is a lot like UNIX. It certainly builds on a third of a century of experience with UNIX. But is that all? Might it not also have benefited from UNIX source code? If that were the case, of course, SCO could profit from it. So they spent a lot of time investigating whether there is any UNIX code in Linux. And, it seems, they found some.

What do you do if you find somebody abusing your intellectual property? The most obvious thing would be to get them to stop. But SCO isn't doing that; instead, so far they have refused to show enough evidence for people to even voluntarily remove it. Instead they require a draconian non-disclosure agreement before they will show the code.^[17]

They appear not to want to have it removed. Chris Sontag, senior vice president of SCO, is on record^[18] as having said: “Once contributed, code cannot be removed.” This is nonsense, of course. Operating systems sources undergo continual development, and that involves replacement of code. SCO’s legal predecessors have fought lawsuits to get people to remove UNIX code from their sources. But this attitude suggests that SCO wants the code to stay there: if Linux contains UNIX source code, they can charge royalties for it.

This would explain a lot of SCO’s actions. But is it true that they want royalties from Linux? And if they do, can they get them?

The effects

SCO’s approach is causing a lot of damage, not the least to SCO.

The most obvious effect of SCO’s actions is FUD: Fear, Uncertainty and Doubt. The letter they sent to their customers appeared to be calculated to cause maximum damage, sparking conspiracy theories that Microsoft was behind the actions, and that the real intention was to kill Linux. These theories gained additional fuel a couple of days later when Microsoft announced that they were licensing SCO technology.^[19]

Other effects are more damaging to SCO. In the complaint they filed against IBM in March, they wrote:

84. Prior to IBM’s involvement, Linux was the software equivalent of a bicycle. UNIX was the software equivalent of a luxury car. To make Linux of necessary quality for use by enterprise customers, it must be re-designed so that Linux also becomes the software equivalent of a luxury car. This re-design is not technologically feasible or even possible at the enterprise level without (1) a high degree of design coordination, (2) access to expensive and sophisticated design and testing equipment; (3) access to UNIX code, methods and concepts; (4) UNIX architectural experience; and (5) a very significant financial investment.

It’s hardly necessary to state that this statement stretches the truth to breaking point. But there are further-reaching implications: long before IBM became involved with Linux, Caldera was active trying to make a “luxury car” out of Linux. This statement implies that they were incapable of doing so. It makes it pretty clear that they no longer intend to be in the Linux business, but it also weakens their case about their technical expertise.

Of course, this is only part of the story. On their web site (<http://www.sco.com/unitedlinux/>), they still advertise their own “enterprise-grade Linux”. This kind of crass contradiction weakens both their case against IBM and their own credibility. Combined with other actions, it also fuels additional conspiracy theories: SCO has shown source code common both to UnixWare and Linux under a non-disclosure agreement, both to press and to a small number of programmers. They claim that the code was stolen from UnixWare and put into Linux.

SCO is not going to great lengths to convince people. Reportedly,^[20] the terms of the non-disclosure agreement

essentially permitted SCO to declare any information it provided to be confidential, regardless of whether the signer already knew it, and which offered no circumstances under which that information could be revealed. Most Linux developers are unable to sign such an NDA, as it easily could prevent them from ever again working on the kernel. Similarly, employees of any company that works with Linux cannot sign such an NDA.

A small number of programmers have seen the code. They have confirmed that there were significant similarities in some areas, similarities which can really only be explained by code copying. SCO has an explanation. In their complaint they claim:

86. It is not possible for Linux to rapidly reach UNIX performance standards for complete enterprise functionality without the misappropriation of UNIX code, methods or concepts to achieve such performance, ...

It would be understandable if someone were to come to the conclusion that SCO themselves inserted the code into Linux: they had both the motivation and the ability.

Could this action improve sales of System V? That’s not very likely. SCO has been selling UNIX all this time, and sales are declining. The big vendors suffering too: that’s why they are moving to Linux. They wouldn’t do that if they thought they could make money out of commercial UNIX.

What about getting royalties for Linux? This seems highly unlikely, but it’s up to a court to decide on that. It’s difficult to charge royalties on code which is not identified. If it is identified, it can be removed. In any case, Linux isn’t the only free UNIX-like operating system. There are the BSDs as well, genuine descendents of UNIX. They had their lawsuits ten years ago, and they have long been settled.

The return of BSD?

Around the time that Linus Torvalds started developing Linux, the CSRG in Berkeley was working to extricate the code that had been written in Berkeley from the AT&T base on which it had been built. That wasn't as ridiculous as it might sound: in over ten years of development, just about all the AT&T code had been replaced. The code which hadn't was relatively simple and easy to replace. Since it was developed at a university, a BSD without AT&T code could be released under a much more liberal license, one that required no royalties.

This project was not completed in Berkeley. Instead, some CSRG members formed a company, Berkeley Software Design (BSDI) to market the new operating system. They had given AT&T's lawyers the opportunity to convince themselves that there was no residual AT&T code in the system, but the lawyers didn't pay much attention until BSDI started advertising their toll-free phone number: 1-800-ITS-UNIX. Then they acted quickly, and within weeks the phone number was changed. But the incident woke the sleeping dogs, and they investigated and subsequently sued.

It's instructive to revisit the results of that action:

- The code in question was derived from AT&T, so the chances were much greater that some original AT&T code might accidentally be left behind. Ultimately they did find some code which was deemed to be derived from the Seventh Edition of Research UNIX. It was certainly nothing that could not easily have been replaced.
- Due to pending litigation, BSDI was effectively unable to market their operating system for two years. During this time, the free systems NetBSD and FreeBSD also appeared on the scene. Their progress was strongly limited by the lawsuit. For years after it was settled, out-of-date press reports kept the illusion alive that there were problems using BSD. Many BSD advocates attribute the meteoric success of Linux to the FUD generated by this lawsuit.
- The case was settled out of court in February 1994. Most conditions of the settlement were kept confidential. FreeBSD and NetBSD did not pay anything, but like BSDI, they had to agree to migrate to 4.4BSD-Lite, a code base which both sides agreed to be "unencumbered".

Could this lawsuit have similar results? Anything's possible, but there are a number of significant differences. In the AT&T case, it was a big company, AT&T, suing a small Colorado startup and involving a couple of user groups. SCO is a small company with a shaky financial foundation^[21] suing the world's largest computer company. In a situation where the ability to finance expensive litigation is important, it's clear that SCO is in a much worse position.

The real mistake

No matter what the outcome of the suit, though, SCO is wrong. The real key to survival is not a few algorithms or lines of source code. Open Source has shown that.

The big lesson of Open Source is that software development works better in the open, where people can discuss things. This benefit far outweighs the commercial value of the source code. UNIX was originally developed in this way, even if access to the source was more restricted than the modern Open Source approach. The problems began in the 1980s, where startup companies tried to outdo each other in a quest for market share. Companies developed their own proprietary extensions to UNIX to lock customers in to their implementation, and they guarded the source code closely, elevating its perceived value far beyond anything that was justifiable. In truth, UNIX source code is no better than Linux source code.

IBM in particular has recognized these mistakes. IBM's corporate culture is about as far from the Open Source ethos as you could find. It's not embracing Linux or releasing the source for its own products out of altruism; it's doing it because it believes it makes commercial sense. In so doing, they *are* gaining an advantage over companies who resist change, a fair advantage. If SCO wants to survive, it needs to understand this lesson. Litigation won't save them.

Greg Lehey <grog@lemis.com> is an independent UNIX consultant living in South Australia. In the course of 30 years in the industry he has performed most jobs, ranging from kernel development to product management, from systems programming to systems administration, from processing satellite images to pro-

gramming gasoline pumps, from the production of CD-ROMs of ported free software to DSP instruction set design. He was the manager of Tandem Computers' European UNIX technical support and has worked in UNIX development at Siemens-Nixdorf and Linux development at IBM. He has been developing BSD code for ten years and is author of the Vinum Volume Manager, a free volume manager for BSD. He is the president of the Australian UNIX User's Group and a member of the FreeBSD core team. He is the author of "Porting UNIX Software" (O'Reilly and Associates, 1995) and "The Complete FreeBSD" (fourth edition, O'Reilly and Associates, 2003). He is maintaining an ongoing analysis of the SCO affair at <http://www.lemis.com/grog/sco.html>.

REFERENCES

1. SCO published this letter at http://www.sco.com/scosource/letter_to_linux_customers.html. As the result of legal action taken in Germany, SCO was obliged to remove it.
2. <http://www.sco.com/scosource/complaint3.06.03.html>
3. <http://www.levenez.com/unix/history.html>
4. *Der große Frust*, iX January 1992, page 75.
5. The systems were AT&T, Interactive, Esix, Microport and Consensus. SCO did not have a System V.4 offering until much later.
6. <http://www.opengroup.org/>
7. <http://www.unix.org/version3/>
8. Compaq is now called Hewlett Packard, but the registration is in the name of Compaq. No Hewlett-Packard product is certified as UNIX 98.
9. <http://www.opengroup.org/openbrand/register/xx.htm>
10. <http://www.opengroup.org/comm/press/who-owns-unix.htm>
11. <http://www.novell.com/news/press/archive/2003/05/pr03033.html>
12. <http://www.novell.com/news/press/archive/2003/05/pr03033.html>:
Put simply, Novell is an ardent supporter of Linux and the open source development community. This support will increase over time.
13. <http://www.levenez.com/unix/history.html>, <http://www.sco.com/scosource/unixtree/unixhistory01.html>. There's also a derivation to Venix 4.2 in 1993, but it never seems to have been distributed.
14. <http://www.opengroup.org/openbrand/register/xu.htm>
15. <http://www.opengroup.org/openbrand/register/xt.htm> <http://www.opengroup.org/products/cert/certprods.htm>
16. <http://www.levenez.com/unix/history.html>, <http://www.sco.com/scosource/unixtree/unixhistory01.html>.
17. <http://www.linuxjournal.com/article.php?sid=6956&mode=thread&order=0>
18. http://www.byte.com/documents/s=8276/byt1055784622054/0616_marshall.html
19. <http://ir.sco.com/ReleaseDetail.cfm?ReleaseID=109360>
20. <http://www.linuxjournal.com/article.php?sid=6956&mode=thread&order=0>
21. <http://ir.sco.com/ReleaseDetail.cfm?ReleaseID=102627>