

Open Source Everywhere

Software is just the beginning ... open source is doing for mass innovation what the assembly line did for mass production. Get ready for the era when collaboration replaces the corporation.

By Thomas Goetz

Cholera is one of those 19th-century ills that, like consumption or gout, at first seems almost quaint, a malady from an age when people suffered from maladies. But in the developing world, the disease is still widespread and can be gruesomely lethal. When cholera strikes an unprepared community, people get violently sick immediately. On day two, severe dehydration sets in. By day seven, half of a village might be dead.

Since cholera kills by driving fluids from the body, the treatment is to pump liquid back in, as fast as possible. The one proven technology, an intravenous saline drip, has a few drawbacks. An easy-to-use, computer-regulated IV can cost \$2,000 - far too expensive to deploy against a large outbreak. Other systems cost as little as 35 cents, but they're too complicated for unskilled caregivers. The result: People die unnecessarily.

"It's a health problem, but it's also a design problem," says Timothy Prestero, a onetime Peace Corps volunteer who cofounded a group called Design That Matters. Leading a team of MIT engineering students, Prestero, who has master's degrees in mechanical and oceanographic engineering, focused on the drip chamber and pinch valve controlling the saline flow rate.

But the team needed more medical expertise. So Prestero turned to ThinkCycle, a Web-based industrial-design project that brings together engineers, designers, academics, and professionals from a variety of disciplines. Soon, some physicians and engineers were pitching in - vetting designs and recommending new paths. Within a few months, Prestero's team had turned the suggestions into an ingenious solution. Taking inspiration from a tool called a rotameter used in chemical engineering, the group crafted a new IV system that's intuitive to use, even for untrained workers. Remarkably, it costs about \$1.25 to manufacture, making it ideal for mass deployment. Prestero is now in talks with a medical devices company; the new IV could be in the field a year from now.

ThinkCycle's collaborative approach is modeled on a method that for more than a decade has been closely associated with software development: open source. It's called that because the collaboration is open to all and the source code is freely shared. Open source harnesses the distributive powers of the Internet, parcels the work out to thousands, and uses their piecemeal work to build a better whole - putting informal networks of volunteer coders in direct competition with big corporations. It works like an ant colony, where the collective intelligence of the network supersedes any single contributor.

Open source, of course, is the magic behind Linux, the operating system that is transforming the software industry. Linux commands a growing share of the server market worldwide and even has Microsoft CEO Steve Ballmer warning of its "competitive challenge for us and for our entire industry." And open source software transcends Linux. Altogether, more than 65,000 collaborative software projects click along at Sourceforge.net, a clearinghouse for the open source community. The success of Linux alone has stunned the business world.

But software is just the beginning. Open source has spread to other disciplines, from the hard sciences to the liberal arts. Biologists have embraced open source methods in genomics and informatics, building massive databases to genetically sequence *E. coli*, yeast, and other workhorses of lab research. NASA has adopted open source principles as part of its Mars mission, calling on volunteer "clickworkers" to identify millions of craters and help draw a map of the Red Planet. There is open source publishing: With Bruce Perens, who helped define open source software in the '90s, Prentice Hall is publishing a series of computer books open to any use, modification, or redistribution, with readers' improvements considered for succeeding editions. There are library efforts like Project Gutenberg, which has already digitized more than 6,000 books, with hundreds of volunteers typing in, page by page, classics from Shakespeare to Stendhal; at the same time, a related project, Distributed Proofreading, deploys legions of copy editors to make sure the Gutenberg texts are correct. There are open source projects in law and religion. There's even an open source cookbook.

In 2003, the method is proving to be as broadly effective - and, yes, as revolutionary - a means of production as the assembly line was a century ago.

In the Beginning

Message-ID:1991Aug25.205708.9541@klaava.helsinki.fi

From:

torvalds@klaava.helsinki.fi (Linus Benedict Torvalds)To: Newsgroups: comp.os.inixSubject: What would you like to see most in minix?Summary: small poll for my new operating systemHello everybody out there using minix-I'm doing a (free) operating system (just a hobby, won't be big and professional like gnu) for 386 (486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhatAny suggestions are welcome, but I won't promise I'll implement them :-)

Linus

Thousands of coders, hackers, and developers answered Linus Torvalds' call - and helped him build a robust system that continues to pick up steam. Yet what's amazing about Linux isn't its success in the market. The revolution is in the method, not the result. Open source involves a broad body of collaborators, typically volunteers, whose every contribution builds on those before. Just as important, the product of this collaboration is freely available to all comers. Of course, there are plenty of things that are collaborative and free but aren't really open source (Amazon.com's book reviews, for instance). And many projects aren't widely collaborative, or are somewhat proprietary, yet still in the spirit of open source (such as the music available from Opsound, an online record label). Not to mention that, as with any term newly in vogue, *open source* is often invoked on tenuous grounds. So think of it as a spectrum or - better still - a rising diagonal line on a graph, where openness lies on one axis and collaboration on the other. The higher an effort registers both concepts, the more fully it can be considered open source.

Of course, for all its novelty, open source isn't new. Dust off your Isaac Newton and you'll recognize the same ideals of sharing scientific methods and results in the late 1600s (dig deeper and you can follow the vein all the way back to Ptolemy, circa AD 150). Or roll up your sleeves and see the same ethic in Amish barn raising, a tradition that dates to the early 18th century. Or read its roots, as many have, in the creation of the Oxford English Dictionary, the 19th-century project where a network of far-flung etymologists built the world's greatest dictionary by mail. Or trace its outline in the Human Genome Project, the distributed gene-mapping effort that began just a year before Torvalds planted the seeds of his OS.

If the ideas behind it are so familiar and simple, why has open source only now become such a powerful force? Two reasons: the rise of the Internet and the excesses of intellectual property. The Internet is open source's great enabler, the communications tool that makes massive decentralized projects possible. Intellectual property, on the other hand, is open source's nemesis: a legal regime that has become so stifling and restrictive that thousands of free-thinking programmers, scientists, designers, engineers, and scholars are desperate to find new ways to create.

We are at a convergent moment, when a philosophy, a strategy, and a technology have aligned to unleash great innovation. Open source is powerful because it's an alternative to the status quo, another way to produce things or solve problems. And in many cases, it's a better way. Better because current methods are not fast enough, not ambitious enough, or don't take advantage of our collective creative potential.

Open source has flourished in software because programming, for all the romance of guerrilla geeks and hacker ethics, is a fairly precise discipline; you're only as good as your code. It's relatively easy to run an open source software project as a meritocracy, a level playing field that encourages participation. But those virtues aren't exclusive to software. Coders, it could be argued, got to open source first only because they were closest to the tool that made it a feasible means of production: the Internet.

The Internet excels at facilitating the exchange of large chunks of information, fast. From distributed computation projects such as SETI@home to file-swapping systems like Grokster and Kazaa, many efforts have

exploited the Internet's knack for networking. Open source does those one better: It's not only peer-to-peer sharing - it's P2P production. With open source, you've got the first real industrial model that stems from the technology itself, rather than simply incorporating it.

There's a reason we love barn raising scenes in movies. They make us feel great. We think, 'Wow! That would be amazing!'" says Yochai Benkler, a law professor at Yale studying the economic impact of open source. "But it doesn't have to be just a romanticized notion of how to live. Now technology allows it. Technology can unleash tremendous human creativity and tremendous productivity. This is basically barn raising through a decentralized communication network."

An Experiment in Open Source

At 37, Jimmy Wales has already established his legacy on the Internet. Seven years ago, Wales, then a futures and options trader on the Chicago Board of Trade, turned the homepages of hobbyists into Bomis.com, an Internet directory that lets visitors catalog related sites into webrings. The result unified the disparate efforts of millions of Internet users. It wasn't open source, a strategy still percolating in software.

But it came close.

Wales wanted something even closer. Long an admirer of Torvalds and free software pioneer Richard Stallman, he had a more deliberate experiment in mind: using volunteer contributors to create an Internet encyclopedia. As in software, perhaps open source could turn consumers into producers.

The first attempt came in 1999 with Nupedia, an encyclopedia project with great ambitions and what proved to be fatally onerous oversight. Aspiring contributors had to apply for access; each article was peer-reviewed and professionally edited. An entry had to make it past seven or eight hurdles before being posted onto the Nupedia site. "After two years and an amazing amount of money," Wales shrugs, "we had 12 articles."

So in 2001, he tried again. Wales and his team eliminated most of Nupedia's barriers to participation and

invented Wikipedia using Wiki, the open source Web-design software. Wikipedia isn't much to look at.

The site resembles cutting-edge Web design circa 1994. But like a lowly Pontiac Sunfire with a modified

computer chip, most of the action is under the hood. A grassroots encyclopedia, Wikipedia has amassed more than

barriers to participation and invented Wikipedia using Wiki, the open source Web-design software. Wikipedia isn't much to look at. The site resembles cutting-edge Web design circa 1994. But like a lowly Pontiac Sunfire with a modified computer chip, most of the action is under the hood. A grassroots encyclopedia, Wikipedia has amassed more than 150,000 entries, using strict open source principles: Anybody can write an article, and anybody else can improve it. Revisions are posted on a Recent Changes page where suggestions are pored over by a dedicated group of Wikipedians. "There's a simple way to tell if it's any good," says Wales. "Find an entry on something you know something about. Odds are it'll hold up pretty well - you'll probably even learn something new."

So what motivates Wikipedia contributors? Pretty much the same things behind any open source project: a dash of altruism, a dose of obsessive compulsiveness, and a good chunk of egotism. It lets users have a hand not just in shaping the debate, but in designing the product. Some are genuinely motivated by the greater good, or find it satisfying to apply their professional knowledge to a broader audience, pro-bono style. And some get

to prove how smart they are.

Not to say mischief-makers don't lurk out there. Wikipedia has banned several ne'er-do-wells from the site, and some areas have been locked down - the front page, for instance, because, Wales says, "people kept putting giant penis pictures on there." But in general, the system works surprisingly well, and the traffic bears that out. This summer, Wikipedia surpassed Britannica.com in daily hits, according to Web traffic monitor Alexa.com. Wikipedia's popularity is all the more extraordinary because, like Linux, it started as a small-scale experiment. But the result challenged Britannica, a 235-year-old institution.

There's some satisfaction in the fact that the technology behind Wikipedia is the same one that's baffled Britannica for years. The old-guard encyclopedia has never figured out how to adapt to the digital era. In 1998, Britannica stopped updating its print version and focused on its CD-ROM, then last year revived the print version. In 1999, it launched a free site online; two years later, switched to a paid version. The struggles aren't unique, but they illustrate how a proprietary model built on traditional notions of intellectual property can be undone by irresistible forces.

Now Wales is thinking big. He wants to square off with Britannica not just online but in print and on CD-ROM. Next year, he hopes to release Wikipedia 1.0, a peer-reviewed and peer-edited compendium of 75,000 entries, available to anyone, for commercial or noncommercial purposes. He's even considered pulling a Red Hat - releasing an affordable paid version - before anybody else does. "Things like textbooks, encyclopedias, dictionaries, reference works - they lend themselves very well to collaboration," says Wales. "In fact, that's how they're done in the proprietary context, too. But it costs Britannica money to pay people to write articles; it costs to edit them. Those are all things we do for free. So how can they compete? Our cost model is just better than theirs."

Harnessing the Means of Innovation

Business gurus have a term for what drove Wikipedia: *innovation!* It's a flaccid buzzword these days, deflated by a decade of leadership seminars and management bibles. But when you look at what's innovative about open

source, think Tom Paine, not Tom Peters.

Open source embodies an ethos as fruitful and resilient as the closed capitalism Bill Gates represents: the spirit of democratic solutions to daunting problems. It's the creed of Emerson, who preached independent initiative and advocated a "creative economy." It's the philosophy of William James, whose pragmatism dictated that "ideals ought to aim at the transformation of reality." It's the science of Frederick Taylor, who proved that distributing work could exponentially boost productivity and replace "suspicious watchfulness" with "mutual confidence." It's the logic of Adam Smith, whose notion of "enlightened self-interest" among workers neatly presages the primary motivation for many open source collaborators.

Finding the roots of open source in Taylor and Smith is especially significant because the approach isn't, as some insist, anticommercial or anticorporate. Rather, it is a return to basic free-market principles. The open

source process fosters competition, creativity, and enterprise. And just as Taylor and Smith provided the intellectual grounding for the revolution in mass production, open source offers the mechanism to mass innovation.

While the assembly line accelerated the pace of production, it also embedded workers more deeply into the corporate manufacturing machine. Indeed, that was the big innovation of the 20th-century factory: The machines, rather than the workers, drove production. With open source, the people are back in charge. Through distributed collaboration, a multitude of workers can tackle a problem, all at once. The speed is even greater - but so is the freedom. It's a cottage industry on Internet time.

Just as the assembly line served the manufacturing economy, open source serves a knowledge-based economy. Facilitating intellectual collaboration is open source's great advantage, but it also makes the method a threat. It's a direct challenge to old-school R&D: a closed system, where innovations are quickly patented and tightly guarded. And it's an explicit reaction to the intellectual property industry, that machine of proprietary creation and idea appropriation that grew up during the past century and out of control in the past 30 years - now often impeding the same efforts it was designed to protect.

Copyright and patents have an admirable purpose: They give creators the right to exploit their creations for a limited time. Then these innovations enter the public domain. If it made sense that copyrights and patents protected products, it makes sense that, in today's economy, they protect ideas and concepts, too.

But the balance and fair-mindedness that made the American system hum like a well-tuned Briggs & Stratton is now clogged up with opportunism. Copyright protections that originally lasted 14 years now drag on for nearly a century, leaving the public domain a barren ground. Particularly since the mid-1990s, when the US Patent and Trademark Office began recognizing business methods, intellectual property has become more than just guarding what you've made. Trademark, copyright, and patents are now offensive weapons. The result often impedes, rather than encourages, innovation. Intellectual property has grown infuriating in its excesses, such as Netflix's recent patenting of something as simple as a subscription model for DVD rentals.

Perversely, this is just how the law wants it. The courts and the patent and trademark offices exist to protect property, be it physical or intellectual - slap on "All Rights Reserved" and reap the rewards. But it's annoyingly difficult to *share* something - to open intellectual property to a wide audience. The conventional legal system simply isn't built to handle "*Some Rights Reserved*."

Open source flips this paradigm around. Now there are dozens of licenses, from Stallman's General Public License to Creative Commons' ShareAlike agreement, that let open products exist in a proprietary world. Under these licenses, to use political scientist Steven Weber's terms, property is something to be distributed rather than protected. The owners are more guardians than guards.

The first and most likely places for open source to flourish are at the extremes of IP. The method can craft better, more open versions of bad business models and inefficient markets. But the imitative projects, the ones that replicate proprietary products using better means, are just the gimmes. In the long term, open source will apply outside IP-dominated industries. Weber suggests corporate R&D as a natural starting point; the oil industry, for instance, could enlist outside chemists to collaborate on better oil refining techniques.

As technology reduces the costs of replication and distribution to nearly nothing, the open source approach could catalyze stagnant sectors of the economy - or, better yet, create new economic sectors. "Open source can propagate to fill all the nooks and crannies that people want it to fill," says Mitch Kapor, who founded Lotus in 1982, cofounded the Electronic Frontier Foundation in 1990, and now heads the Open Source Applications Foundation. (See [Mitch Kapor Reinvents Your Inbox](#).) "In an economy where more and more value is in information - is in the bits, not the atoms, where bits can be copied essentially for free - any time you have that situation, economic schemes that rely on existing models of intellectual property laws for protection are going to do less and less well. If information wants to be free, then that's true everywhere, not just in information technology."

Open Source as a Weapon

A decade ago, Michael Eisen slogged through swamps in Costa Rica studying the

mating behavior of frogs. That's what biologists did, he figured - and if he had to fight off a few leeches along the way, so be it. Now he's all about coding, crafting blocks of genetic data and churning them through his computer. "It's a great time to be a biologist," says Eisen, a computational scientist at Lawrence Berkeley National Laboratory. "*Origin of Species* is the best thing ever written in biology. But you just wish Darwin knew about genomics."

Yet if biology is in a renaissance, there are still relics of a medieval age. Most aggravating to Eisen is the state of scientific publishing. It affronts him. And he wants to destroy it.

His weapon is open source. Unlike Wikipedia's Jimmy Wales, who didn't set out to take down Britannica, Eisen has the publishing community squarely in his sights. Open source, says Eisen, who dabbles in Perl programming, can give rise to a new distribution model for scientific research.

"The guiding principle of science has been that freely available material is more useful; it's more likely to generate better science," Eisen says. But freely available is not the same as free of charge. Science journals, with their historically narrow readerships, often charge thousands for a subscription. One of the biggest disseminators is Elsevier, the science publishing unit of an Anglo-Dutch media conglomerate, which distributes some 1,700 academic journals, from *Advances in Enzyme Regulation* to *Veterinary Parasitology*.

"The whole premise for that model just evaporated with the Internet," Eisen continues. "Technology now makes openness possible; it's maximum openness. The rules of the game have changed, but the system has failed to respond." So Eisen and two colleagues - Stanford biochemist Pat Brown and Harold E. Varmus, a Nobel laureate in medicine and president of the Memorial Sloan-Kettering Cancer Center in New York - have devised an alternative: the Public Library of Science. To Eisen, PLOS is "the optimal system" for publishing scientific research: Where the old model limited access to maximize subscription fees, this library is all about open access - meaning any user can read, download, copy, distribute, print, search, or link to an article.

Instead of charging universities thousands to subscribe to *PLOS Biology* or *PLOS Medicine*, as the library's two peer-reviewed journals will be called, PLOS instead will charge contributing authors a \$1,500 fee to cover costs. (Harvard and other universities typically pay on behalf of their faculties, and the library will waive the fee in special cases.) In addition to hiring editorial staffs, PLOS will distribute the work among a pool of fellow academics. In time, the library will become a storehouse of publicly available scientific research, a resource that, like Linux, will only improve with time. And whereas the old model traditionally had scientists signing away their copyright to the journal, the library will use a license that leaves copyright with the author but allows for unlimited use by third parties, provided credit is given to the author.

"The openness will make the data more useful," says Eisen, pointing to the annual \$57 billion in taxpayer-funded scientific and medical research that isn't available to the public. And then there's the irony that academic institutions get charged for work they did in the first place. "How does it make sense for the universities to give away the copyright to their research and then pay to get access to it again?"

The library follows in the steps of efforts such as BioMed Central, a London-based open-access publisher of online scientific journals. Though BioMed is a valuable repository of all sorts of research - everything from proteomics to psychiatry - it has yet to make a significant dent in the major journals' hold on big research. That's the obstacle PLOS faces, too: convincing scientists that the new publication will have authority. Publishing in *Science* or *Nature* means your paper matters, and it provides a yardstick for tenure and promotion decisions. PLOS has already earned credibility by hiring the former editor in chief of *Cell* and lining up contributors from Harvard, Princeton, Yale, and the National Institutes of Health for the first issues of *PLOS Biology*, set to debut October 13. The team hopes such star power will help put any issues of credibility to rest.

And to those who label PLOS something only a league of anticommercial academics could dream up, Eisen has a swift retort. "It's the ultimate free market - the free market of ideas," he says. "We don't prohibit commercial use, we encourage it. If you want to gather our articles about a particular topic and sell it as a book, great, go ahead. As far as we're concerned, that's a good way of getting information out there. The problem with scientific publishing right now is that it's a monopoly. This is an attack on a specific business model that is not serving science well."

Building a HybridAt any given time, odds are Monsanto is in court over intellectual property. The agriculture and biotech giant is suing rival companies for infringing on technology. It's suing researchers for unlicensed use. And it's suing farmers for stealing seeds - filing 75 or so lawsuits in the past five years alone, most notably against Percy Schmeiser, a Canadian cause célèbre accused of using Monsanto grain without a license; his case goes before the Canadian Supreme Court in January. This constant litigation is a necessary cost of doing business; patents, after all, must be protected, or their value can be lost. But it also reflects a demanding way of doing business, one that's expensive, time-consuming, and - as the Schmeiser case attests - not always good for PR.

For a different approach, consider Cambia, the Center for the Application of Molecular Biology to International Agriculture, a biotech nonprofit based in Australia. Founded in 1994 by Richard Jefferson, a mandolinist turned geneticist from Santa Cruz, California, Cambia has emerged as a force in agriculture technology over the past decade. The group pioneered research into transgenomics, where plants are tweaked using their own genetic stock rather than foreign genes. But Jefferson finds Cambia increasingly hamstrung by the biotech industry's reliance on patents, cross-license agreements, and trade secrets. "So much of what we want to do is all tied up in somebody's intellectual property," he says. "It's a complete sclerotic mess, where nobody has any freedom of movement. Everything that open source has been fighting in software is exactly where we find ourselves now with biotechnology."

So Jefferson tapped open source methods to skirt the restrictive licenses of companies like Monsanto. On a broad scale, Cambia built an exhaustive collaborative database, open to all, of 300,000 patents covering agricultural technologies - an essential resource for researchers navigating through proprietary waters. And as a more precise effort, Cambia is developing a gene-transfer technology that will, Jefferson hopes, work better than the proprietary methods currently available. The group is following Torvalds' model, incubating the core technique before turning it over to a network of users - both nonprofit and corporate - with a liberal licensing arrangement. "Anybody can tweak it, learn from it, twiddle with it," Jefferson says. "We want to invent a better way but bind everybody to share the improvements. Some might call them work-arounds; we call them work-beyonds."

Corporations have been part of the problem for Jefferson, but they're also part of the solution. Open source offers biotech companies a cheaper way to do research. "The corporations have been locked in a zero-sum game," Jefferson says. "It costs them a fortune to buy and lock up a product or a technology. And if they don't, a competitor will get it and they'll have no access to it. So it's a real change in the status quo we're proposing. We're reducing the obstacles for everybody so big companies won't view this as antithetical to their own progress."

Jefferson is onto something. Open source is often framed as an attack on the corporate world at large. But in fact, the open source approach can be a boon for companies. Licensing from a trusted collaborative project saves money and leaves the technology open to further development. By showing corporations that a closed, defensive approach to intellectual property can be less efficient than liberal licensing, Cambia and a few other open source efforts are leading the way to the mainstream.

In this light, where corporations are part of the model, open source suddenly becomes something less marginal and more ingenious. It forces industry to reckon with openness rather than hide behind intellectual property. In driving down the cost of software or encyclopedias or biotechnology, open source is unleashing billions in capital otherwise put to woefully inefficient ends. Just because it's not about making money first doesn't mean it won't make money second (just ask the folks who bought their mansions with Red Hat shares).

"Open source can build around the blockages of the industrial producers of the 20th century," says Yale's Benkler. "It can provide a potential source of knowledge materials from which we can build the culture and economy of the 21st century."

If that sounds melodramatic, consider how far things have come in the past decade. Torvalds' hobbyists have become an army. Britannica's woes are Wikipedia's gains. In genetics and biotech, open source promises a sure path to breakthroughs. These early efforts are mere trial runs for what open source might do out in the world at large. The real test, the real potential, lies not in the margins. It lies in making something new, in finding a better way. Open source isn't just about better software. It's about better everything.