

IBM Began Buying Red Hat 20 Years Ago

How Big Blue became an open-source company.

By *Glyn Moody*

News that [IBM is buying Red Hat](#) is, of course, a significant moment for the world of free software. It's further proof, as if any were needed, that open source has won, and that even the mighty Big Blue must make its obeisance. Admittedly, the company is not quite the behemoth it was back in the 20th century, when “[nobody ever got fired for buying IBM](#)”. But it remains a benchmark for serious, mainstream—and yes, slightly boring—computing. Its acquisition of Red Hat for the not inconsiderable sum of \$34 billion, therefore, proves that selling free stuff is now regarded as a completely normal business model, acknowledged by even the most conservative corporations.

Many interesting analyses have been and will be written about why IBM bought Red Hat, and what it means for open source, Red Hat, Ubuntu, cloud computing, IBM, Microsoft and Amazon, amongst other things. But one aspect of the deal people may have missed is that in an important sense, IBM actually began buying Red Hat 20 years ago. After all, \$34 billion acquisitions do not spring fully formed out of nowhere.



Glyn Moody has been writing about the internet since 1994, and about free software since 1995. In 1997, he wrote the first mainstream feature about GNU/Linux and free software, which appeared in *Wired*. In 2001, his book *Rebel Code: Linux And The Open Source Revolution* was published. Since then, he has written widely about free software and digital rights. He has [a blog](#), and he is active on social media: [@glynmoody](#) on [Twitter](#) or [identi.ca](#), and [+glynmoody](#) on [Google+](#).

OPEN SAUCE

Reaching the point where IBM's management agreed it was the right thing to do required a journey. And, it was a particularly drawn-out and difficult journey, given IBM's starting point not just as the embodiment of traditional proprietary computing, but its very inventor.

Even the longest journey begins with a single step, and for IBM, it was taken on June 22, 1998. On that day, IBM announced it would ship the Apache web server with the IBM WebSphere Application Server, a key component of its WebSphere product family. Moreover, in an unprecedented move for the company, it would offer “commercial, enterprise-level support” for that free software.

When I was writing my book *Rebel Code: inside Linux and the open source revolution* in 2000, I had the good fortune to interview the key IBM employees who made that happen. The events of two years before still were fresh in their minds, and they explained to me why they decided to push IBM toward the bold strategy of adopting free software, which ultimately led to the company buying Red Hat 20 years later.

One of those people was James Barry, who was brought in to look at IBM's lineup in the web server sector. He found a mess there; IBM had around 50 products at the time. During his evaluation of IBM's strategy, he realized the central nature of the web server to all the other products. At that time, IBM's offering was Internet Connection Server, later re-branded to Domino Go. The problem was that IBM's web server held just 0.2% of the market; 90% of web servers came from Netscape (the first internet company, best known for its browser), Microsoft and Apache. Negligible market share meant it was difficult and expensive to find staff who were trained to use IBM's solution. That, in its turn, meant it was hard to sell IBM's WebSphere product line.

Barry, therefore, realized that IBM needed to adopt one of the mainstream web servers. IBM talked about buying Netscape. Had that happened, the history of open source would have been very different. As part of IBM, Netscape probably would not have released its browser code as the free software that became Mozilla. No Mozilla would have meant no Firefox, with all the knock-on effects that implies. But for

OPEN SAUCE

various reasons, the idea of buying Netscape didn't work out. Since Microsoft was too expensive to acquire, that left only one possibility: Apache.

For Barry, coming to that realization was easy. The hard part was convincing the rest of IBM that it was the right thing to do. He tried twice, unsuccessfully, to get his proposal adopted. Barry succeeded on the third occasion, in part because he teamed up with someone else at IBM who had independently come to the conclusion that Apache was the way forward for the company.

Shan Yen-Ping was working on IBM's e-business strategy in 1998 and, like Barry, realized that the web server was key in this space. Ditching IBM's own software in favor of open source was likely to be a traumatic experience for the company's engineers, who had invested so much in their own code. Shan's idea to request his senior developers to analyze Apache in detail proved key to winning their support. Shan says that when they started to dig deep into the code, they were surprised by the elegance of the architecture. As engineers, they had to admit that the open-source project was producing high-quality software. To cement that view, Shan asked Brian Behlendorf, one of the creators and leaders of the Apache project, to come in and talk with IBM's top web server architects. They too were impressed by him and his team's work. With the quality of Apache established, it was easier to win over IBM's developers for the move.

Shortly after the announcement that IBM would be adopting Apache as its web server, the company took another small but key step toward embracing open source more widely. It involved the **Jikes Java compiler** that had been written by two of IBM's researchers: Philippe Charles and Dave Shields. After a binary version of the program for GNU/Linux was released in July 1998, Shields started receiving requests for the source code. For IBM to provide access to the underlying code was unprecedented, but Shields said he would try to persuade his bosses that it would be a good move for the company.

A Jikes user suggested he should talk to Brian Behlendorf, who put him in touch with James Barry. IBM's recent adoption of Apache paved the way for Shield's own

OPEN SAUCE

efforts to release the company's code as open source. Shields wrote his proposal in August 1998, and it was accepted in September. The hardest part was not convincing management, but drawing up an open-source license. Shields said this involved “research attorneys, the attorneys at the software division who dealt with Java, the trademark attorneys, patents attorneys, contract attorneys”. Everyone involved was aware that they were writing IBM's first open-source license, so getting it right was vital. In fact, the original Jikes license of December 1998 was later generalized into the IBM Public License in June 1999. It was a key moment, because it made releasing more IBM code as open source much easier, smoothing the way for the company's continuing march into the world of free software.

Barry described IBM as being like “a big elephant: very, very difficult to move an inch, but if you point the elephant toward the right direction and get it moving, it's also very difficult to stop it.” The final nudge that set IBM moving inexorably toward the embrace of open source occurred on January, 10, 2000, when the company announced that it would make all of its server platforms “Linux-friendly”, including the S/390 mainframe, the AS/400 minicomputer and the RS/6000 workstation. IBM was supporting GNU/Linux across its entire hardware range—a massive vote of confidence in freely available software written by a distributed community of coders.

The man who was appointed at the time was what amounted to a Linux Tsar for the company, [Irving Wladawsky-Berger](#), said that there were three main strands to that historic decision. One was simply that GNU/Linux was a platform with a significant market share in the UNIX sector. Another was the early use of GNU/Linux by the supercomputing community something that eventually led to *every single one of the world's top 500 supercomputers* running some form of Linux today.

The third strand of thinking within IBM is perhaps the most interesting. Wladawsky-Berger pointed out how the rise of TCP/IP as the de facto standard for networking had made interconnection easy, and powered the rise of the internet and its astonishing expansion. People within IBM realized that GNU/Linux could do the same for application development. As he told me back in 2000:

OPEN SAUCE

The whole notion separating application development from the underlying deployment platform has been a Holy Grail of the industry because it would all of the sudden unshackle the application developers from worrying about all that plumbing. I think with Linux we now have the best opportunity to do that. The fact that it's not owned by any one company, and that it's open source, is a huge part of what enables us to do that. If the answer had been, well, IBM has invented a new operating system, let's get everybody in the world to adopt it, you can imagine how far that would go with our competitors.

Far from inventing a “new operating system”, with its purchase of Red Hat, IBM has now fully embraced the only one that matters any more—GNU/Linux. In doing so, it confirms Wladawsky-Berger's prescient analysis and completes that fascinating journey the company began all those years ago. ■

Send comments or feedback
via <https://www.linuxjournal.com/contact>
or email ljeditor@linuxjournal.com.