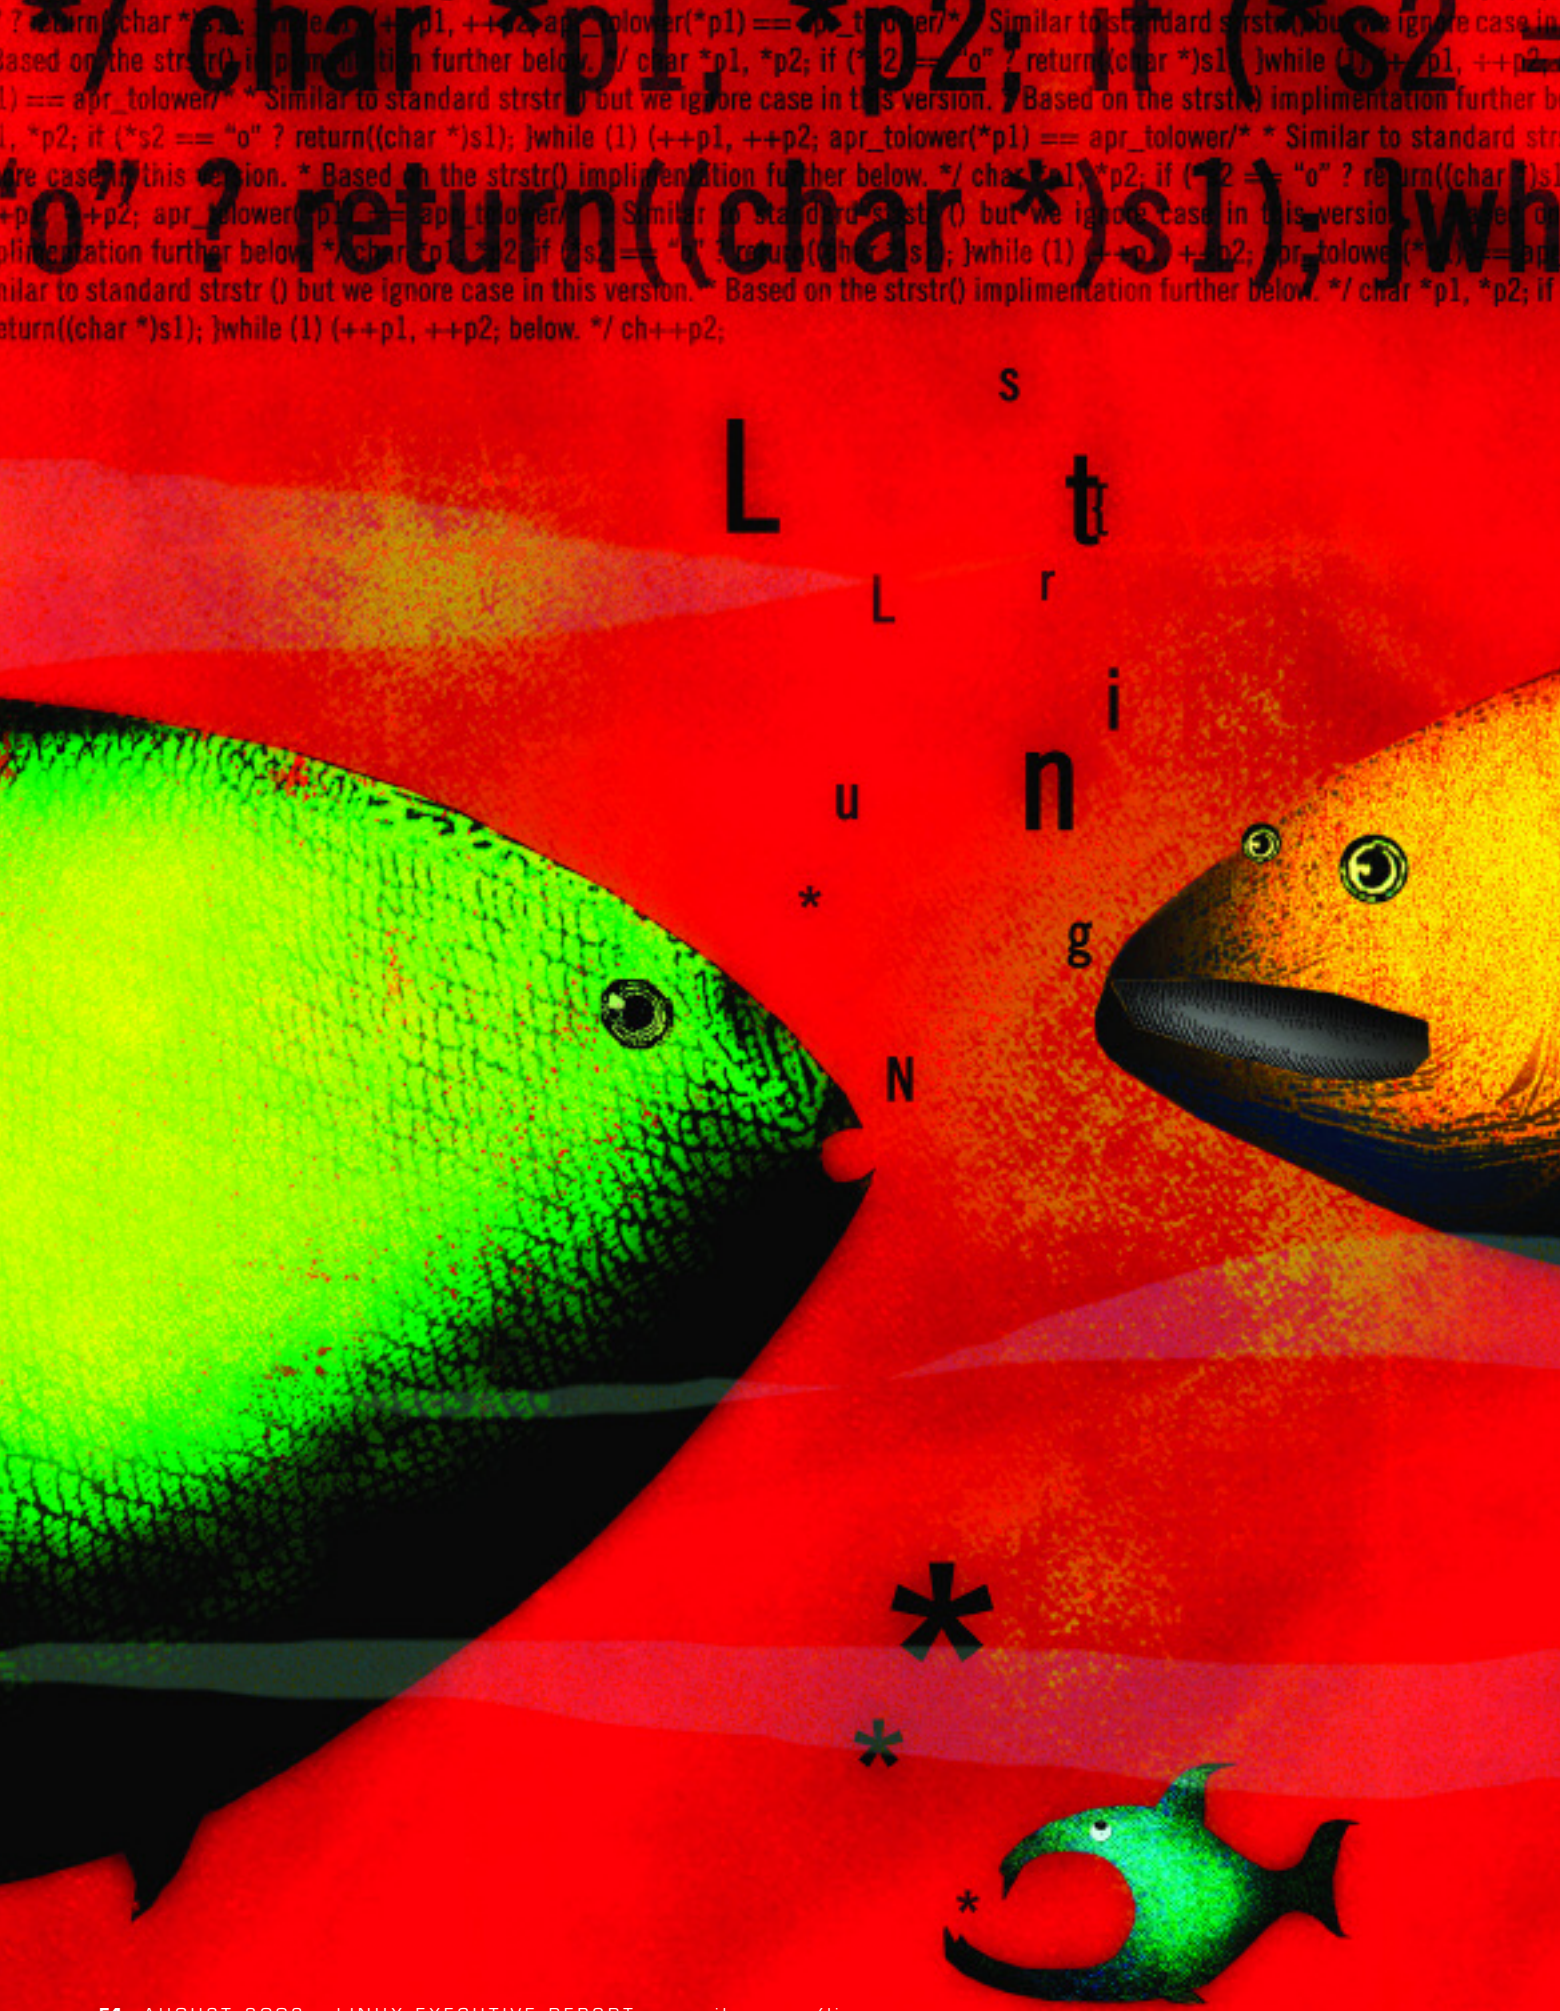
L
i
n
u
x

Past, Present and Future

Three executives discuss the whys and wherefores of IBM's approach to open source beyond Linux

BY JIM UTSLER

The open-source community has come a long way since Linus Torvalds developed and then openly shared Linux*. Now, more and more applications, development tools and other technologies are being put up on the block, ready for developers of all stripes—from a guy working at home to an army of corporate programmers—to add to them.

Linux, of course, is still considered the flagship of all open-source projects, but others are making their way into mainstream corporate IT environments. *Linux Executive Report* (LER) recently discussed this trend with three IBM executives—Scott Handy, vice president, Worldwide Linux and Open Source; John Palfreyman, director of Open Source, Grid and Virtualization services; and Jeff Smith, vice president, Open Source and Linux Middleware—to get their perspectives.

ILLUSTRATION BY BRIAN STAUFFER

LER: Why are customers now becoming more interested in open source beyond Linux?

JS: I think there's been quite a bit of success with Linux, and people are now saying, "Hey, I wonder if this can apply to other things." So one reason is the success of Linux, which leads users to look at other open-source options. Another thing is the new business model that surrounds open source, where you pay less for the license and more for the support and subscription. Those models are getting people's attention, and they're investigating whether the results of that can offer them new or better value than they've been able to get with whatever they have been using in the past. A third thing is that customers are realizing more and more the value of building solutions around open standards, and open source projects are a good way to get broad proliferation of open standards.

SH: That's right, because when they're using an open source technology, they're using something that everybody else is also using so they can get interoperability and adherence to a particular standard. Of course, most business are working in a mixed environment, with customers having both open source technology and maybe some private source as well.

JP: As Jeff points out, some customers consider open source because they're looking for cost savings; they're obviously attracted by the fact that the software does not attract any license fee. However, they should do so with both eyes open. Yes, license fees don't exist, but that doesn't mean there aren't costs involved in the project. So we need to make sure they consider all of the TCO (total cost of ownership) implications of what they're doing, and the other benefits that will accrue from making the right choices.

LER: What are some of the challenges customers are facing when adopting open source beyond Linux?

JP: Let's look at that a bit. We advise customers to look at their business needs first and not to adopt open source just for open source's sake. And the challenges here can be several fold. They need to determine how to create the appropriate combination of open-source and proprietary software, how to assess the maturity of the open-source software-how to make sure that the community support is there and it's not just from one vendor who's actually marketing their product in an open-source license-and to understand that there are certain things that need to be taken care of with open source, such as configuration management and control. With a traditional piece of software, the vendor will take care of this. This isn't the case with open source. Customers have to make sure they really understand the full and holistic implications of using open source in their business.

JS: Integration is also an issue. Customers who bring components of open source back into their environment find that they go through more work to integrate those things than they might with a vendor who built some of that integration into their commercial offerings, as in the case with the WebSphere* family. For example, WebSphere Application Server (WAS) uses the open source Apache HTTP server as a key component of an integrated software solution implementation.

SH: I think the challenges when implementing open source are the same as with any technology deployment. Customers should have all of the same types of plans in place for supporting that technology that they would for any other private source technology. Once customers see it in that light, they're much more comfortable with the list of things they need to do to deal with in an open-source environment. They often can just use the same checklist that they've used when deploying any new technology, all the way down from how to get it approved on their standard deployment list to then rolling it out into production.

"Apache is important because it is an open-source community that proliferates the Web, open Web standards, and Web-applications standards."

—Jeff Smith, vice president, Open Source and Linux Middleware, IBM

LER: Why is IBM so involved in the open-source space?

SH: Based on our experience with Linux, we figured out that some very useful things are developed with open source and that, often, it's actually in both the customers' and IBM's best interest to use open source in the model and the mechanism for how we deliver innovation to customers. From our perspective, it's certainly better than having multiple vendors doing duplicate development work to create more or less the same product. It's a waste of resources. But by sharing the exact same code base, we, as a vendor, can take our valuable resources and use them for developing value add on top of that standard, whether that be Linux, whether that be an open-source J2EE app server, whether that be Eclipse. This philosophy applies to every open-source project that we've backed. We can share the development work for the open source and open standards and then provide more value to the customer by implementing a unique value-add on top of it.

JP: We're also seeing customer interest in open source. And we see a market there because customers see competitive advantage if they are able to combine open source and private source, and we believe that IBM is well positioned to help them make objective, business-led decisions to maximize their competitive differentiation from harnessing this combination.

JS: In keeping with that thought, we've found that commercial ecosystems form around cores of open-source community innovation that include the business opportunities that John mentions. It's one of the reasons that open-source projects have gotten the attention and the investment that they have. Linux is a prime example of that. People are selling billions of dollars of hardware under Linux and middleware on top of Linux and services surrounding Linux. The starting premise is that there's a core of community-driven innovation that an ecosystem builds around. IBM's goals regarding that innovation are to participate in the ecosystem in a way that is mutually beneficial to our customers, to IBM and to the ecosystem itself.

"What you want to do is use the capabilities of open source where appropriate, which probably will be in a mixed open-source/private-source environment. If you can get a little in, get a little in, if you can get a lot in—the more the merrier."

—Scott Handy, vice president, Worldwide Linux and Open Source, IBM

LER: How has IBM's experience with Linux helped it with other open-source projects?

JS: First of all, we've learned about what does and doesn't work when you're working with open-source communities. We've become knowledgeable about the communities themselves: how they function, how to contribute to them and how to earn their respect, which is important for you to be able to get value from them. We have shown that this idea of commercial ecosystems forming around these open-source community innovations can be successful, which has lead us to consider more of them. And we've shown that customers are interested in this stuff, that alternative solutions, especially those that offer more choice to do things, are inherently interesting to customers. So not only do we see value from our business strategy perspective, but our customers to a certain degree are also voting for us to participate in these things.

SH: We've also learned that to make an open-source project successful, it has to be repeatable. Let's take a quick look at Linux: We learned that there was a certain set of things we knew we had to do to be successful in the community, including becoming a major contributor, not only for credibility with the customers, but also for credibility within the community. And that can only be done through real technical contributions of innovation into the open source community. That allows us to advance the technology so that it meets the needs of our customers, so if we have customer requirements, we know how to feed those requirements into the technology. It also allows us to provide support offerings for the technology because we're familiar with the code base and have a mechanism to support our customers.

JP: I'd add to that. Because we're extending our services portfolio to embrace open source beyond Linux, we're extending the skill sets of our professionals who already have experience with Linux. We're also used to interacting with a product that is open source, so this gives us a sound knowledge base to build on. However, it would be foolish of us to expect that the entire spectrum of open source will evolve in exactly the same way as Linux. Open source is very multifaceted, with penetration into areas such as tooling and middleware and embryonic penetration into the application arena.

LER: What are some of the major areas of IBM investment in open source?

JP: Our main areas of investment largely align with what the Software Group has been doing. We're clearly looking to make sure that we include open source in our services portfolio, to build on the investment in proactive community development that has been made in the Software Group, with things like Eclipse and Apache and the other communities that we're very involved in.

SH: Right now, there are several areas that I'm focused on. One involves the application-server space, which we've done by adopting Apache as our default HTTP server in WebSphere and continue to do by working within the ecosystem development around Apache Geronimo. We also created the Eclipse project when we open sourced our technology, which involved 40 million dollars worth of development, back in November 2001. Since then, Eclipse has been a phenomenally successful project, with more than 800 tools developed for it. And building on Eclipse in the area of the rich-client platform (RCP) that was in Eclipse 3.0. That was really for extending Eclipse beyond app-development tools to support building client-side apps. So instead of developing tools, we're developing apps. This is a very important open-source based project for us because this really is nothing short of re-creating what we did on the server, of having a single programming model that is multi-OS and based on open standards to provide customers with a more cost-effective way to deploy apps that are multi-vendor and multi-OS. We also felt that there were a lot of benefits to using middleware on the server, so we injected client-side middleware into the strategy and we did that by developing our own framework on top of Eclipse 3.0, which is the IBM Workplace* managed client.

We noticed that customers had to deal with storage management that was vendor unique in all cases. For example, IBM was solving it one way with our

storage, EMC another way-all the various data-storage vendors had their unique way of doing storage management. All customers wanted was to have a common storage-management solution. So IBM took a leadership position and open sourced some of our technology as well as contributions from other companies to launch the Aperi project, which is building an open-source implementation of the open standard SMI (Storage Management Initiative) backed by SNIA. And moving beyond software, we then started thinking we could extend the same open-source software concept to hardware. We're now licensing the specifications for the Power architecture* to interested companies, and there are currently around 40 companies that have joined power.org. That's one hardware implementation. Another one we did was with blade.org. So we're already sharing the specifications for the blade and chassis architecture with other companies, and now there are more than 600 companies that are involved in blade.org. Another area that I'm focused on involves grid. We started getting involved in open source grid project back in the year 2000, and there was a whole specification that was being created called open grid services architecture-OGSA-and we've implemented that grid architecture specification into our own products, such as WebSphere and Tivoli* software.

JS: We also still have hundreds of people writing code for Linux fulltime, contributing into the Linux open-source community, porting Linux to IBM hardware platforms or writing middleware for it, so the open-source involvement with IBM on Linux is fundamentally the biggest. The second biggest is probably Apache. I don't want to repeat too much of what Scott has already said, but Apache is important because it is an open-source community that proliferates the Web, open Web standards, and Web-applications standards and Web-services standards that are fundamental to our software strategy. As an example, we announced the availability late last fall of WebSphere Application Server Community Edition (WAS CE), which is based on an open source J2EE application server project from Apache called Geronimo, and we even bought a company that had a commercial offering based on this technology that formed the basis of our WAS-CE offering. We donated some embedded Java* database technology a couple of years ago, and that has now graduated to a full-blown database project inside Apache called Derby. We continue to offer that as a commercial offering called CloudScape as well as a bundled version with the WAS CE offering. And then, of course, there's Eclipse, which may be just as big as Apache ... at least from an IBM participation perspective.

JP: The other important thing to stress is that we decided to join up on our open-source plans to make sure that there are no overlaps and gaps in our approach. So that's something that's really exciting-actually having something truly coordinated between the three different divisions-software, systems and services-and getting a lot of benefit from bouncing things off each other and moving these activities forward together.

LER: What advice would you give to IBM customers considering adopting open source?

"Because we're extending our services portfolio to embrace open source beyond Linux, we're extending the skill sets of our professionals who already have experience with Linux."

—John Palfreyman,
director of Open Source,
Grid and Virtualization services

JS: First of all, if you're used to buying software in a traditional way, where you buy it up front from a particular vendor and it's their own technology, then there's never been any question about where you're going to get support or how much support you're going to get. But what customers are finding with open-source projects is that there's no guarantee there's any support at all; that you actually have to pick where you're going to get the support from, because it's often offered by more than one vendor. Or it could be that the only support available is the Web interface into the community, in which case it's not going to be as predicable as you might like if you're an enterprise customer. Support options may exist, but you have to deliberately seek them out.

JP: The advice I would give is stand back and consider your real business rationale before acting. Open source must never be viewed as the universal answer to all questions, but if used in the right combination with traditional software it can yield significant business benefit, allowing the customer to innovate. If the answer is "We don't need open source yet because it's really not mature enough in the area that we want to do business," that's a really good answer. Increasingly, however, that won't be the answer.

SH: In keeping with that, what you want to do is use the capabilities of open source where appropriate, which probably will be in a mixed open-source/private-source environment. If you can get a little open source into a project, get a little in, if you can get a lot in—the more the merrier. But don't set these arbitrary rules that everything has to be 100-percent open source when 30 percent might be enough to get some advantage. Take your success and expand on it. In short, just bite off what you can chew and go from there.

Jim Utsler, *MSP TechMedia senior writer, has been covering the technology beat for nearly a decade. Jim can be reached at jutsler@msptechmedia.com.*