

# The State of Open Source

Contributed by Sukh

The night that divides the two days of tutorials from the three-day conference at the fifth annual O'Reilly Open Source Convention is reserved for the States of the Union addresses. Luminaries from the open source communities of Perl, Python, PHP, MySQL, Apache, and Linux each spoke for just under a half hour to present their take on the current state of their technology and where it is headed.

## The Perl State of the Onion

Larry Wall kicked off the evening with an entertaining look at what he called the Universal architectural diagram. Looking much like the court markings on one side of a tennis court, the diagram was filled in a variety of ways during a lightning session of political, cultural, religious, and geek references. In one version, the base of the rectangle was Parrot and the three vertical rectangles that sat on top of it were filled with languages that would run on this single engine. The left rectangle held Perl 5, the middle Perl 6, and the right contained a large list of languages that included PHP, Ruby, Python, BASIC, Scheme, COBOL, and Java. Wall explained that currently only BASIC is supported and wondered if this larger goal is really impossible.

To illustrate his point, he refactored the image into a widget--the canonical impossible object. He addressed the criticism of how slow Perl 6 development is by adorning the prongs of the widgets with the labels "fast", "good", and "cheap". He argued that the traditional wisdom is that you can pick two but he has advocated that you can pick only one, and that he has chosen "good". Although the basic design of Perl 6 is mostly done, Wall cautioned that he will eventually have to get a real job with health insurance.

Originally the goal had been to have one implementation of Perl. This has changed to a plan to have three: the current Perl 5, which is based on C, and an implementation of Perl 5 and Perl 6, each based on Parrot. The Perl 5/Parrot stack is designed to help people make the transition from Perl 5 to Perl 6 more easily. This will enable you to translate modules one at a time and run some under Perl 5 and some under Perl 6, all running on top of Parrot. Wall also noted that Perl 5 runs faster on Parrot.

Perl 5 over Parrot has been named Ponie to avoid confusion with the traditional Perl 5 implementation. Ponie will be compatible with Perl 5.10. Wall's final piece of news was that Fotango will sponsor Arthur Bergman to work on company time on the implementation of Ponie.

## The Python State of the Snake

Guido van Rossum began his address with slides that showed Python's growth over the past few years and outlined plans for the upcoming 2.3 release. The timetable for the Python 2.3 release has been accelerated because of its inclusion in the next version of Mac OS X. The Panther release of the Apple operating system is scheduled to ship by the end of this year. When it was noted that the Python 2.3 beta was included in the Panther beta given to developers last month, the Python release schedule was compressed so that the release candidate is scheduled for July 27 and the final release is planned for August 3.

One feature that didn't make it into Python 2.3 is conditional expressions. van Rossum proposed a syntax for a new conditional. After much debate the syntax that was voted on was (if C:x else:y). It passed a vote but not by much. According to van Rossum, "The voting instructions were kind of an IQ test in themselves." The margin was small and the debate continued about feature bloat on the one hand and the benefits of the feature on the other. The conclusion was to follow the conservative path: "when in doubt, don't change anything".

Next, van Rossum looked ahead to Python 3.0, the point at which backwards compatibility may be broken. He explained "We're throwing away a lot of the cruft that Python has accumulated." He offered many examples. Strings were objects with no methods. Python 3.0 will replace the string module with string methods. As an aside, van Rossum noted that Python may run on Parrot. There are two methods that generate collections of integers; xrange() generates a sequence of integers while range() generates a list of integers. One is more efficient and one is easier to understand. The goal is to get rid of xrange() and let range() perform both functions. The issue of 8 bit versus Unicode strings will be addressed. The ability to use map and filter to process lists will disappear, being replaced by list comprehensions. Exceptions will be expressed by classes and not by strings.

As for the Python 3 schedule, van Rossum said "it will not be any time soon. Until the Python Foundation has money to fund me, Python 3 will remain vaporware." He hopes that the community as a whole takes over with tiny nudges from him, which he can fit into his current work schedule.

## The PHP State of the Elephant

In the past year and a half there have been seven PHP releases. Shane Caraveo said that the big release in the past year was December's 4.2.3. There are still threading problems with Apache 2.0 that are being addressed. More recently, PHP is experimenting with a command-line interface for shell scripting. Streams I/O has been implemented, and PHP

now works on Mac OS X. After almost two years, PEAR 1.0 has been released. PEAR will provide some of the capabilities of CPAN by providing a repository of script modules. PECL is a separate and complementary distribution for C Extensions.

The PHP 5.0 release features revamped internals and real object orientation. 5.0's object support is much better than 4.0. Caraveo is careful to note that this is not going to take away from functional use of PHP. Not everything will be turned into objects. Caraveo said that "the implementation of objects will include constructors and destructors--a lot of the things you'd want to see in objects if you actually want to use objects." He continued, "The heart and sole of PHP is its ease of use. In the future PHP will advance but still stay with its core."

The new release will introduce exceptions and an improved streams API. Also, SQLite will be bundled to provide an SQL interface over a flat file system. Caraveo explained, "There have been rumors that we're removing MySQL support from PHP. That's not true. We removed the MySQL client library from the distribution because most people already had the client library."

The MySQL State of the Dolphin

David Axmark and Monty Widenius of MySQL introduced the new features in the current stable and development versions of MySQL. (Axmark showed a picture of Monty's daughter, My, claiming that Widenius named the software after her. Widenius has a younger child, Max, who was apparently the inspiration for MySQL-MAX.)

The conference T-shirt comes from MySQL. The front says "MySQL." The back simply has the word "Transactions" with a checked box. That's much the state of MySQL--adding long-requested features. Both Axmark and Widenius explained that subselects are working (and fast) and stored procedures are working, as per the SQL-99 specification. Axmark commented that only IBM's DB2 followed the procedure language guidelines in the specification.

Interestingly, MySQL AB supports customers on at least four versions of its product. They have support contracts going back to version 3, but the development is very clearly pushing toward the 5.0 release. Widenius promises it by the end of the year, with 5.1 to follow shortly.

MySQL has expanded greatly, even in the past few weeks. Though the home office is in Sweden, most developers work from home. In one sense, Axmark joked, they could claim that they have around 70 offices. The core development team has tripled, from around 9 developers to nearly 30 now. Additionally, MySQL now employs three full-time documenters, including Paul DuBois.

The Apache State of the Feather

Greg Stein, Chair of the Apache Software Foundation, presented a state of the Apache Foundation and not a state of the Apache web server. The Apache Software Foundation was founded to provide organizational, legal, and financial support to its now more than 100 active projects. Stein said that a goal is to "enable code to survive any individual developer." Three people may start a project and the project might grow to ten committers while three disappear. What rights and control should that original group of three retain? The ASF also provides a long-lasting viable home for software and, as a corporation, can get software grants from Sun, IBM, and other companies, and can provide legal agreements to transfer that code.

Stein summarized Roy Fielding's presentation on the Apache Vision, saying, "Real innovation occurs when contributors are able to implement their own ideas at their own pace, and for their own purposes. The idea is to provide a framework where people can contribute without the worries." The Apache Software license is intended to make the software usable and modifiable by the largest number of people possible. Stein articulated the underlying principle as being, "Trust the community to do the right thing. Do not try and force contributions. No viral qualities or heavy limits on modifications." He referred to Jon Stevens' warning that "A license can ruin a perfectly good piece of software."

As for the future, Stein urged the audience members to start by joining a mailing list, submitting good patches, and working their way up to committer or member. Participation is important and this is not limited to programmers. There is a need for non-coding members who perform public relations, write docs, or help to provide infrastructure.

The Linux State of the Penguin

IBM's Theodore Ts'o presented his "What's Up with Linux" address. Work on 2.5 began in November 2001 with feature freeze almost a year later at Halloween 2002. The 2.5 close down is in progress as more bugs are being addressed. Ts'o believes that 2.6 will come out in the next couple of months. Much of the remainder of his talk looked ahead to what is expected in the 2.6 release, such as new architecture support, including x64-64, ppc64, user mode Linux that runs Linux inside of Linux, h8300, and uClinux.

The list of improvements to Linux is extensive. The module loader has been reimplemented. There is a new Kernel Build System. Now make with no arguments builds the appropriate compressed kernel and appropriate modules for the appropriate architecture. Ts'o noted that although Eric Raymond's CML2 and Keith Owens kbuild2.5 didn't make it into

the tree, many ideas from kbuild 2.5 were included. The block I/O layer has been improved. Many of the efficiency gains are because the buffer cache is no longer used as the primary interface by block device driver. The VM subsystem has gone through massive changes to make 2.5 more stable under heavy loads than 2.4 is today.

Kernel preemption has been added to 2.5 so animations will be less jerky as will be the playback of MP3s. In addition, writing CDs can now use DMA to rip audio and to write CDs. Ts'o joked that these additions to 2.6 "help Linux become the RIAA's worst nightmare."

The laundry list of improvements continued. In the network there is now more support for bridging firewall code. IPSEC has been added and uses the KAME tools, not FreeSWAN, though people are porting the FreeSWAN userland tools to the KAME API. The scheduler has been improved--this is Ingo Molnar's O(1) scheduler. Ts'o pointed to problems with huge numbers of threads spawned in some Java applications. He explained that "Threads are an abstraction. So 32 threads on a 16-processor system makes sense, but not the 20,000 threads that Java sometimes required." In 2.6 this problem was solved by supporting larger numbers of threads. 2.6 is a much more scalable kernel than 2.5. There is also support for the new Posix-compatible threading library, NPTL. Filesystem improvements include POSIX ACLs and Extended attributes. Newly supported filesystems include NFSv4, JFS, XFS, and CIFS. Many other filesystems have been improved.

Ts'o concluded that although much has changed, much of the Linux kernel is just being refined. The keys are better scalability and better support for larger machines. The exciting things are happening in user space.