This penguin is the Linux official mascot, chosen by Linus Torvalds to represent his image of the operating system he created. Others question whether "a fat penguin really embodies the grace of Linux."

Fermilab for Linux

by Sharon Butler, Office of Public Affairs

The 20-something Finnish computer whiz who invented the Linux operating system e-mailed Fermilab last week on learning that the Laboratory would not only be using the free software for its on-line and off-line analyses but had decided to support the system (see www.fnal.gov/cd/CDN/CDN-jan98/cdn-1.html).

"I'm wringing my hands maniacally, and laughing in a very disturbed manner. 'Whahahahhaaa, I'm taking over the world, yes, yes, YES!'" Linus Torvalds joked.

On a more serious note, though, Torvalds said he was not surprised. The operating system has always had its base in technical and university research centers.

Just two years remain before Run II. With its explosion in the number of particle collisions, Run II places unprecedented demands on Fermilab's computing facilities. Experimenters will be collecting at least 20 times more data than in Run I, requiring a potentially costly upgrade in the Laboratory's data-processing capabilities.

What to do?

Intel's Pentium-Pro 200-MHz chip, introduced to the market in 1995, opened up the possibility of using commodity PCs instead of the more specialized UNIX workstations used in Run I—at a much lower price.

Add to the PCs the Linux operating system, which is free of maintenance costs, and the savings, Fermilab realized, would become even more significant.

Fermilab began toying with the idea of switching to the more cost-effective PCs fitted with Linux software, but the idea took off only when G.P. Yeh, a Fermilab physicist in the Computing Division working in the CDF collaboration, got involved.

The crucial question, Yeh said, was "Would it work?"

As a test, Computing Division staff—including Don Holmgren, Don Petravick, Ron Rechenmacher, Jim Fromm, Connie Sieh and Ken Stox—built a small cluster of PCs for off-line data analyses. Yeh had his colleagues in CDF's Taiwan group transfer the collaboration's computing code, which runs on Silicon Graphics and IBM UNIX workstations, to the PCs and make them run using Linux. To their surprise, the task took a mere three weeks—not a year or more as expected. "It was really simple," said Yeh, "because Linux is just a generic UNIX."

Meanwhile, Holmgren and Andy Beretvas, of the Computing Division, ported many of the standard high-energy software programs used in off-line analyses, with equal success.

Yeh and his group went further. They suggested using PCs with Linux operating systems for Level III triggers, computer farms that process the on-line data from a particle detector. In Run I, Fermilab had used high-end Silicon Graphics Symmetric Multiprocessors because of the large network bandwidth required. Advances in switch technology and networking in general have since made the PC option feasible. CDF's Massachusetts Institute of Technology group, which is responsible for upgrading the collaboration's Level III triggers, tried the new idea, and Petravick, Holmgren and Rechenmacher again got involved, figuring out how to interconnect the computers in the high-speed network. No surprise this time. The idea easily worked.

Yeh laughed, "Now people are saying 'It's so obvious, we should use PCs and Linux.'"

Computing Division staff are now supporting Linux. And scientists in the E871 and E815 experiments, the Theoretical Physics Group and the Sloan Digital Sky Survey collaboration are using the PC/Linux configuration.

And so, Yeh said, "we are quietly changing the way we compute."

Torvalds is also pleased. Before he got "completely side-tracked into computers," he wrote, he was "one of the math-physics geeks" and "very interested in particle physics.... I still feel kind of proud that I'm involved even if it is in a fairly distant manner." ■

Photo by Reidar Hahn

(Left to right) Ron Rechenmacher, of the Computing Division, Kevin McFarland, a physicist from the Massachusetts Institute of Technology, and G.P. Yeh, a Fermilab physicist, examine Run I data using the new off-line test farm.

Linux for Smarties

In four years, CEO Robert Young and his partner, Marc Ewing, took Red Hat, Inc. from a tiny start-up to the leading global supplier of Linux, the legendary open source computer operating system invented by young Finnish computer whiz Linus Torvalds in 1992.

Torvalds, a university student, had the revolutionary idea of creating a clone of the Unix operating system and making it available free of charge, along with its underlying source code, to any user who wanted it. In 1995, Young and Ewing founded Red Hat, Inc. to market software packages that make Linux more user friendly. Red Hat put together Linux packages with third-party applications, documentation and technical support and sold them for about $50 apiece. Sales were brisk.

So brisk in fact, that Red Hat soon became the leading supplier of Linux-based operating system supplies. When the company made an initial public offering of its stock on August 11, 1999, the IPO was a major media event. The price of the stock soared from $14 to $52 in a single day.

Now, readers can get an inside account of Red Hat's brief but hectic life in *Under the Radar, How Red Hat Changed the Software Business and Took Microsoft by Surprise,* by Young and co-author Wendy Goldman Rom. (Coriolis Press, $27.50). The book tells how Ewing and Young went "from selling Linux out of our homes (to avoid getting real jobs) with few ambitions for great financial success, to being fought over by the world's two largest investment banks during our IPO."

It was, by all accounts a wild ride, and Fermilab was along for parts of it. In fact, Fermilab is on the scene as the curtain rises in *Under the Radar*'s opening chapter, "Inside the Tent":

There was a blip on the screen, something new in the field.

At first barely visible. It had appeared slowly and almost imperceptibly. Indeed, at first it had been difficult to see there was anything at all. No, this was not some dramatic sighting, no alien mother ship suddenly blazing its way across the screen.

The first time that engineers at silicon giant Intel Corp. had the first inkling of change was when scientific labs across the country began demanding that it port its "math libraries" to a new operating system.

For one, Dr. Yeh, a Taiwanese scientist at Midwest-based Fermilab, had made such a plea in early 1998. Fermilab, the federally funded atom-smashing think tank overseen by the U.S. Department of Energy, was a mecca for the world's top nuclear [sic] *physicists. It had quietly added a new flavor of system software to its roster of those driving the lab's network of computers.*

Such sites were known in the computer industry as "early adopters," technically savvy users that often were the first to install leading-edge products before the market had fully accepted them. One of the critical benefits of the new software that Fermilab had installed was that it was almost crash proof, and—even more importantly—scientists could freely tinker with its source code, the guts of any piece of software.

This was not the norm in the Microsoft-dominated software industry. Source code was like a secret chamber that few were allowed to enter. By keeping this code to themselves, software companies kept control of their customers, dictated technological change, and ensured continual revenue streams. With the source code kept secret and inaccessible, customers were locked into continual operating system upgrades dictated by the supplier. Likewise application software creators depended on the internal workings of the operating system and were often put at a disadvantage by the suppliers' secrecy.

"We need your math libraries to run under Linux," a number of Fermilab scientists repeatedly told Intel.

As we know, the blip on the screen has become a worldwide computing phenomenon, with an estimated 20 to 30 million Linux users, several hundred of them at Fermilab.

"Dr. Yeh," that "early adopter" from the federally funded atom-smashing think tank somewhere in the Midwest, was of course Fermilab physicist G.P. Yeh, who was indeed a leader in the laboratory's exploration, testing and ultimate adoption of Linux for Run II computing applications. Today, the Computing Division supports Linux, and its Fermilab applications continue to grow.

So it's fitting that the inscription on the flyleaf of Yeh's personal copy of *Under the Radar* should read "Thanks for your and Fermilab's help! Cheers, Bob."

—Judy Jackson

A Geek In Paradise

Figure 1. Fermi Campus

A trip to see the particle accelerator at Fermilabs by a self-professed geek.

by Jon "maddog" Hall

I had been to Fermilab only the year before, but when the invitation came from Dan Yocum to meet at Fermilab's facility outside Chicago, how could I refuse? I am a geek at heart.

Figure 2. Dr. G.P. Yeh (third from the left) and Linux supporters: Ruediger Oertel from SuSE, Fermilab System Administrator, G.P. Yeh, Stefan Traby from Quant-X, Larry Augustin from VA Linux Systems, Norman Jacobowitz, Linus Torvalds, Dan Yocum, maddog and Matthew Cunningham

Fermilab is short for "Fermi National Accelerator Laboratory", located in Batavia, Illinois. It occupies a parcel of land about three miles on each side (see Figure 1), and houses several accelerator rings which generate (in a *very* concentrated space) amounts of power greater than those found in the sun or any other place in the galaxy, much less on the face of the earth. They use these fantastic amounts of power to collide various particles at extremely high speed in the search for the basic building blocks of the universe.

Figure 3. No, it's not a set from *Star Wars, Episode 1*, it's the Fermi main building.

In ancient days, various philosophers stated that we would eventually find the "smallest particle", and for a while this was considered to be the atom. In the relatively recent days of discovering nuclear energy, it was recognized that the smallest particle was *not* the atom, but made up of various other parts such as protons, neutrons and electrons. (Students of physics, please have mercy on me as I try to explain this in words that most readers will understand.) During the last quarter of a century, more and more physicists began to believe there were even smaller particles making up the protons, called quarks and gluons. Quarks (having nothing to do with a resident of *Deep Space Nine*) are thought to have six different types, and in 1994 the last of these Quarks, the "top quark", was discovered at Fermilab. Unfortunately, the top quark exists for only a very short (10^{-24} seconds) period of time, so it is very hard to collect data on it, particularly when it is seen only six times in a given year of running the accelerator. Therefore, Fermilab decided to increase the size and power of its accelerator, so it could see anywhere from 20 to 300 times the number of quarks. Unfortunately, this would take anywhere from 20 to 300 times the amount of power and generate 20 to 300 times the amount of raw data to be seen by the collectors, meaning 1,000,000MB of data would be generated every second. Yes, that is one million megabytes of data per second.

Of course, storing that much data would be very difficult, but fortunately Fermilab had determined they would be able to filter the information and store a smaller subset of it (only 18 to 100MB of data per second) for later analysis. To do this, they would have to increase the power of their computing systems significantly, and their former model of using expensive workstations in a workstation farm would not have been affordable. Enter Linux.

Last year, when people from Red Hat Software and I visited Fermilab while attending Spring Comdex, I was lucky enough to meet G. P. Yeh, a big fan of Linux and one of the physicists who discovered the top quark. He was kind enough to take us on a short tour of the Fermilab facilities and explain the role of Linux within Fermilab. He explained they investigated Linux and proved that inexpensive PCs running Linux could do the job more than adequately for a price they could afford. They estimated they would need about 2,000 CPUs working together.

Figure 4. Collider Rings

This year, when Dan Yocum heard that Linus Torvalds was speaking at Spring Comdex, he enlisted my help in convincing Linus to make a separate trip to Fermilab to speak to the physicists and their families. This did not take much convincing, since Linus has an interest in math, physics and science.

We met at the hotel where Linus was staying, and with a small group of Linux supporters (see Figure 2), drove to Fermilab. It is quite interesting to approach Fermilab, since the land around the accelerator is flat, with only the main building (see Figure 3) rising up from the ground to any height. It would definitely be a great scene for a science fiction movie.

We parked the car, went inside and met Dr. G. P. Yeh (who everyone calls "G.P.").

G.P. took us on an extended tour, beginning with the top floor of the main building, looking out over the collider rings. "As far as you can see in every direction is Fermilab", G.P. said. It was an impressive sight. He then took us to see the collider detectors (see Figure 4)—"It weighs only 100 tons and cost about 100 million dollars." Finally, we visited the computer room, where the Linux Farms were going to be placed (see Figures 5 and 6). Fermilab calls their systems "Farms" rather than Beowulf systems. They have master machines that delegate the work to many slave processors, connected by high-speed networking and switches. They are not planning on buying the 2000 CPUs until very close to the time they need them. After all, prices keep dropping and capabilities keep increasing, so why not wait until the last moment to get the best "bang for the buck"?

After the tour was over, we went to the main auditorium where Linus gave his talk. For those of you who have heard Linus give a speech, you know he does not like to talk with prepared slides, but instead gives a short prepared talk, then answers questions. This night was no different, other than the topic and complexity of the questions. It was obvious from the questions asked that the audience had more of a computer science bent than other, more general audiences. Questions regarding symmetric multi-processing and the reality of distributing interrupts over multiple CPUs entered the air.

Figure 5.Computer Room–stacks of Linux boxes

Figure 6. Linux Farms: Larry Augustin from VA Linux Systems, Dan Yocum from Fermilab

After a significant amount of time answering questions and signing autographs, our little troupe went to the home of Jeff Gerhardt to enjoy pizza and "refreshments". We were greeted by smoke rolling out of the front door, reminding everyone it is best to take the pizza out of the box before warming it in the oven. When the smoke died down, some interesting home brew made its way to the front, and everyone enjoyed the pizza and brew (see Figures 7 and 8).

Figure 7. Party Time: Linus on left by lamp, G.P. Yeh in far chair, Stefan Traby in far right

Figure 8. Jeff Gerhardt's hospitality (and kitchen) were enjoyed by all.

Figure 9. G.P. Yeh shows map to Linus Torvalds and Stefan Traby. World Domination begins at Fermilab!

I love this type of computing where people push the envelope of what the human mind can conceive, and I thank the government of the United States for helping to fund such a quest.

Jon "maddog" Hall is Senior Leader of Digital UNIX Base Product Marketing, Digital Equipment Corporation. He is Executive Director of Linux International.

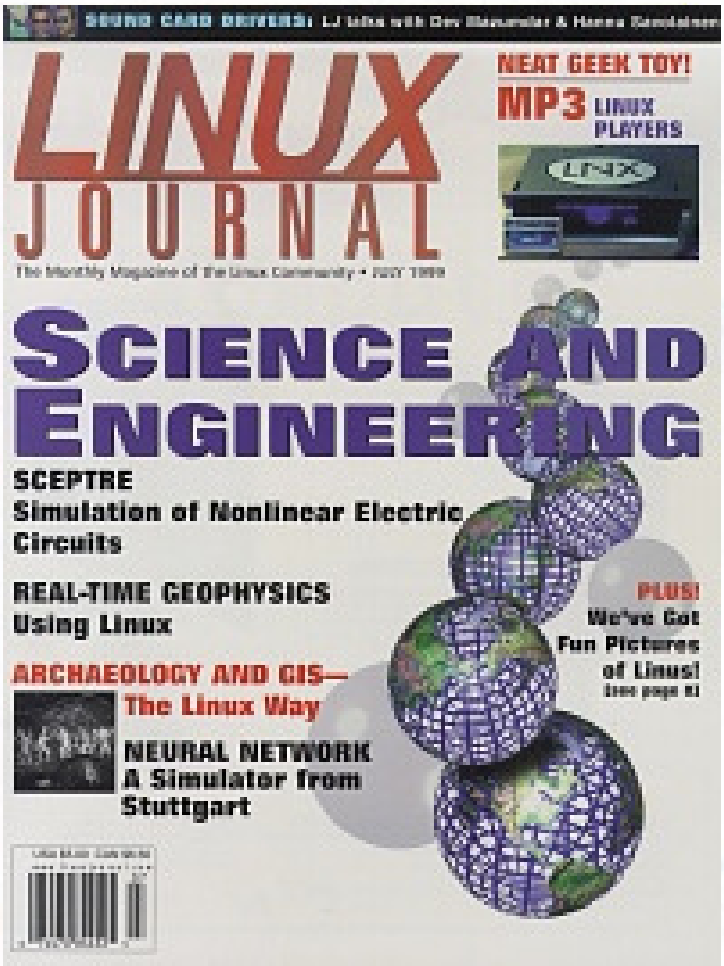

Computing Power

10^9 MegaByte
PetaByte Data Handling

Fermi Lab ⇔ world

- 300 Gbps network
- Grid Computing
- Collaboration

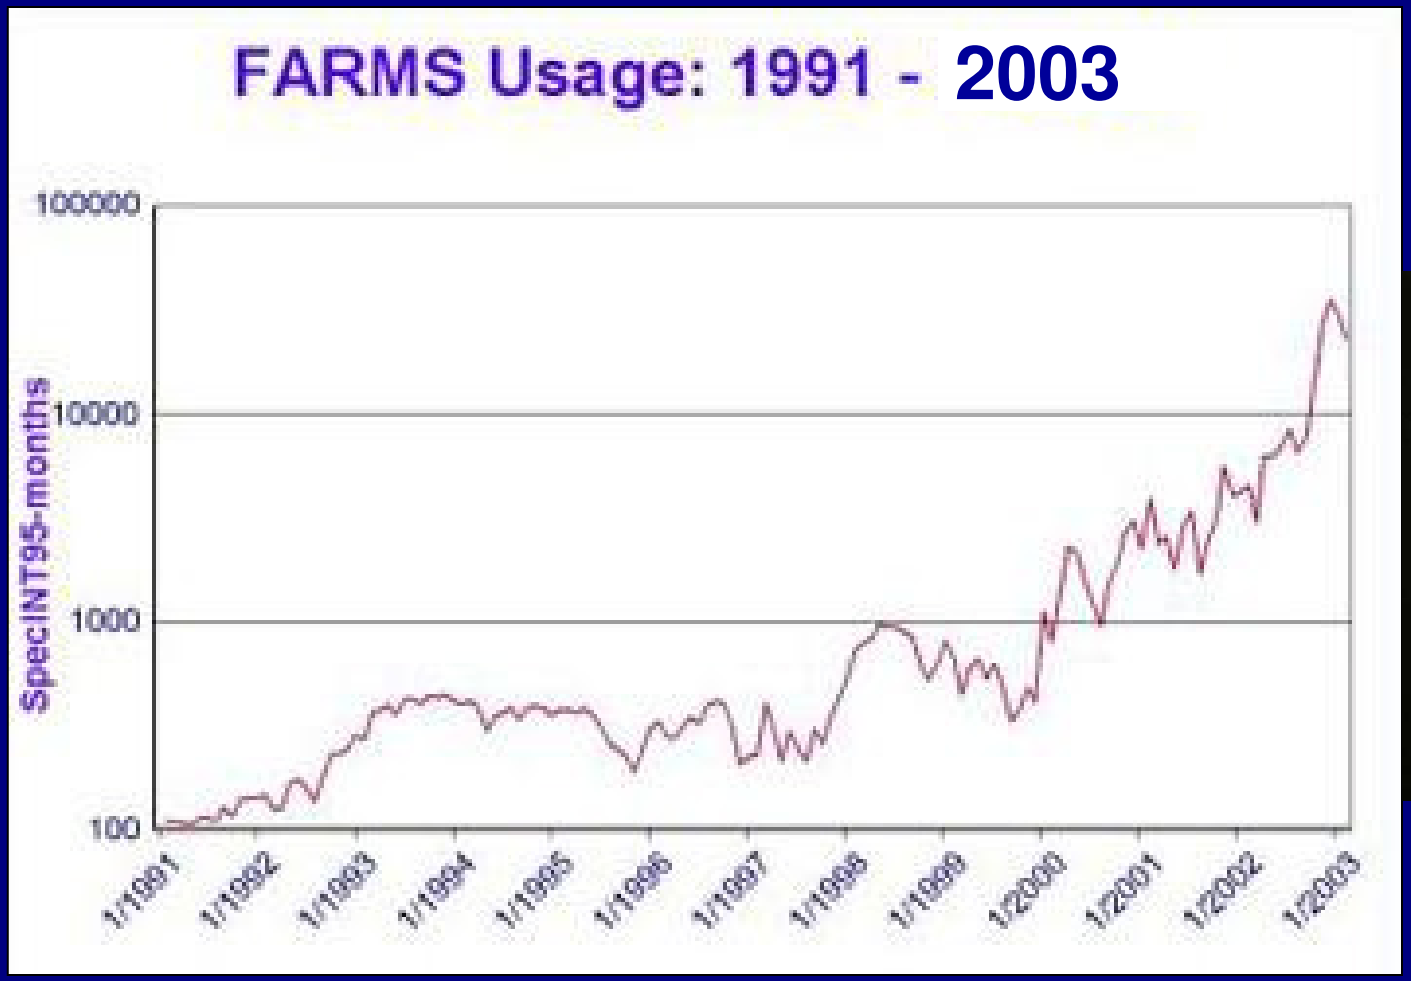

FermiNews

Fermi National Accelerator Laboratory

Volume 19 | FRIDAY, FEBRUARY 23, 1996 | Number 4

The CDF Remote Control Room

Getting in on the Action, from Afar

Photo by REIDAR HAHN

G.P. Yeh, seated at a demonstration unit of the CDF remote control room in the lobby of Wilson Hall. The monitors on the upper right display cross-sections of the detector. Monitors on the left show the "head-on" view of the colliding region in the detector, and a "lego plot" of the secondary particle energy. Yeh is filmed by a small video camera mounted at about eye-level; he's looking into the display from those cameras on the monitors below. Two graduate students are seated behind Yeh.

by Leila Belkora, Office of Public Affairs

An invisible hand draws a white circle against a black background, almost filling the screen of a computer monitor. The circle represents the outer edge of the CDF detector at Fermilab. A spray of green, blue, and red lines blooms rapidly from the center of the circle, revealing the trajectories of particles stemming from the latest collision of a proton and an antiproton in the Tevatron accelerator. One green line arcs toward the upper right of the screen, two more curl around to the lower left. Faster than you can say data visualization, the invisible hand draws a yellow box around the green line on the right, the path of the particle with the highest calculated momentum. The screen goes black again. The image of another proton-antiproton collision at CDF—an ordinary event, as these collisions go, or a rare one that will send up a flag to physicists on the experiment—is due on the screen in less than ten seconds.

continued on page 8

Inside