

## IDE: A Brief Cultural History

# Kernel Development

If you've been confused about the state of the Linux kernel IDE driver in recent months, you're not alone. IDE has been a mess for a long time, with some particularly violent explosions, and some notable successes in the recent past. **BY ZACK BROWN**

IDE, short for "Integrated Drive Electronics", is the communication protocol used along the ribbon of wires connecting your hard disk to the rest of the system.

It's not just another simple driver, or a simple standard. Originally designed by Compaq around 1986, the effort to standardize the IDE interface began in the late 1980's.

Unfortunately this mix of standards pertaining to the same type of hardware, coupled with dizzying variations in manufacturer compliance to those standards, make for a support nightmare. Mark Lord handled IDE driver maintenance for years in the mid 1990's.

It's in the nature of the IDE driver, or anything related to data storage, always to be a prime suspect when a user experiences file corruption. The IDE drivers often took the blame for any filesystems corruption even though Mark and the other developers often found that the real problem lay with other parts of the kernel or even hardware failures – an awkward situation.

### Hot Summer

In the Summer of 1998 the situation boiled over. The kernel was struggling to get to 2.2, and there were still various reports of silent disk corruption with DMA support enabled in the IDE driver. At this time SMP support was also

enabled by default. Having both on by default made it impossible to narrow down where the real faults lay.

Common ground could not be found in deciding which should concede their position. When Linus released 2.1.111, he decided to force the issue. With a one-line patch of his own, he disabled DMA entirely, making it not even a configuration option. Users simply could not use DMA without patching the kernel by hand, even though he admitted that the file corruption that folks had reported was almost certainly due to broken or noncompliant hardware.

### New Start

He had overplayed his hand. Mark abandoned maintainership of the IDE driver, saying that there were no clear cases of data corruption that could be traced to the IDE driver, and that the driver did the most sensible and sane things by default. During this time, Andre Hedrick was one of the more active IDE developers, though some personality conflicts were already starting to show themselves, foreshadowing some of the troubles to come in later years.

Andre officially took over maintenance of the IDE driver and by that September's 2.1.122 release, was hard at work with a bunch of other folks including Mark, getting things ready for 2.2.

Over the next few years, the IDE driver grew to support many variations of disk hardware, in various stages of compliance. Everyone agreed that it would be completely unacceptable to have a situation in which a user could lose data due to poor Linux support of their hard drive. To keep abreast of the particularities of the drives coming out on the market, Andre connected himself with various companies and standards bodies, signing NDAs that would allow him to gain access to privileged information about upcoming hardware, but not to share that information with other developers.

In spite of support for a broad array of IDE hardware, the basic complexity of

the code remained; and this, coupled with the enforced industry secrecy, made it extremely difficult for developers to make any changes to the IDE driver. Certain things that looked like obvious cleanups, turned out to break things in hard-to-see ways. The published specifications left much to be desired, especially for newcomers, with industrial secrecy obfuscating much of the necessary information needed to make patches work correctly.

By February of 2002 the IDE code was very sophisticated, but very, very messy, with even top developers finding their code unreadable. Andre and perhaps a few others seemed to understand the protocols and various compliance issues, but no one else seemed able to get near it.

Because of this, few patches were accepted, with the notable exception Marcin Dalecki, who began feeding in-

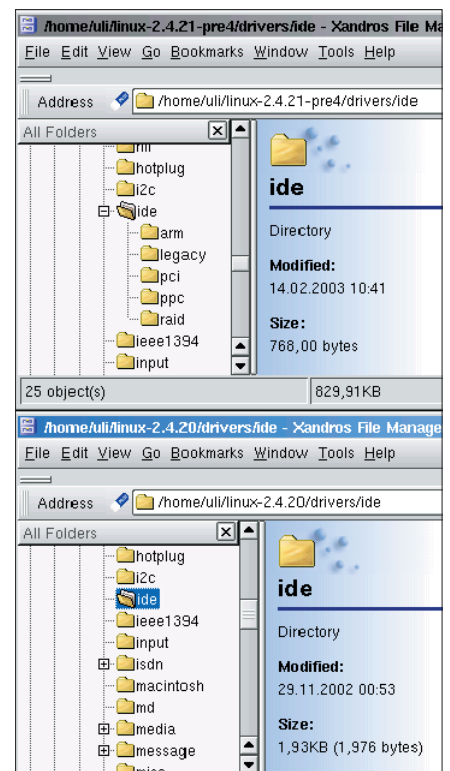


Figure 1: The work within the kernel for the IDE driver is showing benefits after a slow confusing start

vasive ‘cleanup’ patches to Andre at an ever-increasing rate. But Andre, perhaps stressed out by years of IDE nightmares, made some strange accusations, that the aggressive speed of Marcin’s patches were because of a hidden agenda at the venture capital firm employing Marcin. Marcin just laughed at this, but conflict continued to mount.

The issue of Marcin’s patches became more and more controversial, with some developers telling him to quiet down and work better with Andre. At one point Linus summed up his take on situation by saying, “Guys, you need to realize that Marcin is NOT the bad guy here. The problem is that Andre cannot take any level of criticism, and in the five years or so that he has been maintainer I have yet to see a *single person* who has been able to work together with him.”

By this time Andre was already threatening to leave off IDE development in the 2.5 tree, and Linus was accusing Andre of lying about the behavior of patches. Everyone involved found the situation intolerable.

## Clean sweep

When the dust had settled, Linus had made Marcin IDE maintainer for 2.5, with a specific mandate from to clean up the code. Throughout the Spring of 2002 Marcin pursued a campaign of gutting the IDE driver. In retrospect, we can point out ways in which he might have done things differently, but at the time it was very difficult to see how best to proceed. The IDE driver contained years and years of compacted filth, that had to be shaken out. For several months, Marcin shook. And the more he shook, the more stuff fell out.

Not everyone was in favor of this. A lot of developers felt that Marcin was not adequately replacing the ugly code with code that was better designed. In spite of the criticisms, Marcin kept it up, along with folks like Bartłomiej Żolnierkiewicz who was a key contributor. And Linus made it a point to respond to criticism of Marcin’s work, justifying the harshness of the overhaul.

But the 2.5 IDE rewrites couldn’t help but make developers very nervous in general, and a number of folks refused to have anything to do with 2.5 until IDE

settled down, it was clear that the IDE driver was becoming less and less stable. This had a noticeable impact on the amount of work being done on the kernel as a whole, with 2.5 development slowing down.

There was very little choice at that time. Marcin had to keep going forward with his cleanup, but that left other developers up a stump. Finally in July 2002, Jens Axboe decided to forward port the old 2.4 IDE code back up to 2.5. He was careful to say that this was not intended to replace Marcin’s code, just to present a temporary alternative to developers who felt stymied by the current situation.

A lot of people were very happy to see this and kernel development picked up again. It was suggested that Marcin should have done something like this from the start.

In August of 2002 Marcin threw in the towel. Bart had gone off to do his own independent IDE fork; and according to Linus, Marcin had just gotten fed up with constant criticism. With Marcin gone, and Andre disqualified because of personality conflicts with Linus, there were some back-room discussions about who would take on the responsibilities of IDE maintainer. Marcin’s 2.5 IDE work was far from finished, and the code was in a horrendous state.

Eventually Alan Cox agreed to act as maintainer, and to be a buffer between Andre and Linus. Alan is believed by many to be originally from Krypton, but even he has limits. Nevertheless, he managed to shoe-horn the massive IDE project into his schedule. Jens also took on part of that responsibility, accepting IDE patches from Andre and others, and passing them up to Linus.

## Rocket speed

At this point, things really started to happen. Andre had a plan for long-term restructuring of the IDE driver into something that sane people could work on and stay sane. He showed this to Alan, and Alan took off with it, moving so fast on design and implementation that no one could keep up, not even Andre. Andre ran alongside as fast as he could, offering his expertise in questions of protocol compliance and particular hardware exceptions to the specifications,

while Alan did the transformation that no one had been able to achieve in all the years of Linux kernel development.

When the smoke had cleared, there was still a lot to be done, but the basic nature of IDE development had been changed forever. The bulk of this recent work was not actually done in the 2.5 tree as one would expect for such a delicate operation. Some of the tricky stuff did take place in 2.5 first, but it was backported to 2.4 as soon as possible, and the main body of the IDE rewrite was done as patches against Marcelo Tosatti’s 2.4 tree. The main reason for this was that the 2.4 IDE driver still had all those years of muck in it.

It seemed to Alan and Andre that the best way to deal with that was to work on the 2.4 changes from the start, rather than come up with a 2.5 solution and back-port to 2.4 later. In December, Marcelo accepted Alan’s big IDE patch, and flipped the switch from the old driver to the new. Shouts of jubilation were heard across the world.

## Bright future

IDE development is ongoing, and a lot of work remains, including the forward port to the 2.5 tree. Even in 2.4, there are still breakages to be fixed, more code to be cleaned and made maintainable. Some things that had previously worked, were left in a somewhat broken state after the restructuring.

For example, for a while IDE could not be compiled as a module, and had to be integrated directly into the kernel binary itself. Alan, refusing to distract himself with these surface-level issues, refused even to look at patches dealing with these and other problems, until the deeper structural changes were firmly in place. But eventually, the ‘little’ things did start to get taken care of.

Whether Alan will continue as IDE maintainer is an open question. Perhaps Jens will take over at some point, or even Andre. New drives continue to come out on the market, and continue to find new ways to creatively interpret the standards. Meanwhile there are hints that a new kid on the block, called serial ATA, may help alleviate some of the difficulties of the old hardware, and the old specifications. ■