

ZACK'S KERNEL NEWS

Cleaning House

Andrew Morton has posted a patch to schedule FUTEX_FD for removal by June 2007, saying that it suffers from “unfixable races”. This condemnation was also the standard used against DevFS in the old days, though a clear explanation of why the races were unfixable was hard to come by. It is doubtful FUTEX_FD will be revived. It has very few (or no) users, and according to its author, Rusty Russell, it was always something of a hack and a special case. Once it’s removed, Rusty says, it should be possible to yet further simplify the futex code. In that case, it’s possible that Andrew will want to accelerate the removal schedule.

Meanwhile, Adrian Bunk continues the largely thankless task of removing old, unmaintained, and broken code from the kernel. He posted a patch recently to remove the VIDEO_ZR36120 driver, which has been marked as broken for three years. As Mauro Carvalho

Chehab put it, folks have said at various times that they would fix the driver, but then nothing happens.

Adrian has also posted patches to remove the MAC89x0, ATARI_BIONET, ATARI_PAMSNET, and SKMC drivers. These have also been marked as broken for three years (since 2.6.0), and have not been fixed in that time. In response to these patches, Geert Uytterhoeven said that he had a patch lying around by Matthias Urlichs, which might fix the MAC89x0 driver. However, as Matthias said, this would need to be checked by someone before it should be included.

Adrian has also posted patches to remove the FB_CYBER FB_VIRGE, FB_RETINAZ3, FB_ATARI, FB_SUN3, and FB_PM3 drivers, also marked as broken for three years. This time, James Simons stepped up to try to maintain these drivers if he could get hold of the relevant hardware. Michael Schmitz pointed out that Geert Uytterhoeven had revived the FB_ATARI driver with a preliminary patch, which Michael then updated and submitted.

Adrian has also posted a patch to remove the SKMC driver, also broken for three years. No one stepped up to claim it.

the project active would be a better maintainer. Josh’s initial release had about 20 contributions from Alexander Viro, Adam DiCarlo, and Pavel Roskin, as well as Josh. A later 0.2 release had about the same number of patches (this time, mainly bugfixes), from Josh and Pavel, as well as Christopher Li, Damien Lespiau, Nicolas Kaiser, and Oleg Nesterov.

Status of CREDITS

Jarek Poplawski recently pointed out that the CREDITS file was very much out of date, missing a lot of people who had made massive contributions over the years. Linus Torvalds explained his current philosophy of the CREDITS file – it was useful five years ago, when there was no version control in use for kernel development. At that time, the only way to keep track of who had made contributions was to list them all in a file. With revision control, contributors are tracked in the revision system itself, in much finer detail than the CREDITS file ever could. Linus regards the CREDITS file as a legacy – still useful for identifying early contributors, but not by any means a canonical source of information, especially for recent work.

Status of sysctl

I recently said in this column that sysctl had been deprecated and would be removed. At the time, I said that unless some big user-space application was found to depend on sysctl, it would probably have a short life.

Well, Eric W. Biederman, who originally scheduled sysctl for removal, has learned of some legitimate sysctl users, with some interesting results. First of all, sysctl will no longer be marked as deprecated, and will be compiled into all kernels by default. Thus, as Eric says, “we get all of the advantages of a fast interface for anything that needs it. Currently sys_sysctl is about 5x faster than /proc/sys, for the same string data.”

However, sysctl is still on the chopping block. With all these existing users, it can’t just be deprecated and removed,

Sparse Maintainership

Josh Triplett has taken over maintainership of Sparse from Linus Torvalds and put out the first official release, version 0.1. Linus never actually put out any numbered releases himself. The Sparse tool parses C code (ANSI C and a number of GCC extensions), analyzing things like which locks a bit of code acquires or releases, or what address space a pointer points to. These bits of information don’t necessarily tell whether code is broken or not, but they help a human reader analyze the code by eye.

Linus started working on Sparse in 2003, providing a set of features sufficient to his needs, after which his maintainership became much less active. When it became clear that other users wanted to add more features, he suggested that someone with more interest in keeping

The Linux kernel mailing list comprises the core of Linux development activities. Traffic volumes are immense, often reaching ten thousand messages in a given week, and keeping up to date with the entire scope of development is a virtually impossible task for one person. One of the few brave souls to take on this task is Zack Brown. Our regular monthly column keeps you abreast of the latest discussions and decisions, selected and summarized by Zack. Zack has been publishing a weekly online digest, the Kernel Traffic newsletter for over five years now. Even reading Kernel Traffic alone can be a time consuming task. Linux Magazine now provides you with the quintessence of Linux Kernel activities, straight from the horse’s mouth.



but the original reasons for removing it (Andrew Morton has called the sysctl code “goop”) are all still valid. So, Eric plans to work with distributions to identify existing users and help them migrate away from sysctl, until it once again seems feasible to deprecate and remove the code. His current time estimate is a year or two, until it would make sense to consider removing the code again.

Support for Apple Motion Sensor

Stelian Pop and Michael Hanselmann have implemented a driver for the Apple Motion Sensor (AMS), found on the 2005 revision of the Apple PowerBooks and iBooks. Various other folks also contributed numerous fixes. The patches had once upon a time lived in Andrew Morton’s -mm tree, but he’d removed them after receiving conflicting patches from the developers. As Stelian put it, things have settled down since then, and is ready to return to the -mm tree for broad testing.

Relocatable bzImage

Having to load the kernel into a one-megabyte location has placed a size burden on Linux development over the years. Compression has helped to mitigate that burden, but always there has been the desire to break through that barrier, to have larger kernels that support more features. However, the problem is fairly thorny. In particular, code to suspend and resume a running system tends to be very demanding of any kernel relocation code.

Eric Biederman had implemented some patches to support a relocatable kernel a while ago, and Vivek Goyal at IBM has been testing them for some time. Vivek now feels the time is right to get Eric’s patches into Andrew Morton’s -mm tree for broader testing. Because Eric is now involved in other projects, Vivek has forward-ported the patches to the latest kernel, and added his own fixes and improvements.

A lot of kernel hackers have jumped onto this new effort, and in addition to

the relocation features themselves, valuable cleanup is going into these typically dense and messy portions of the kernel.

Support for the Siemens SX1 Phone

Vladimir Ananiev submitted a patch to add support for the Siemens SX1 phone. Tony Lindgren and Pavel Machek were enthusiastic about the patch, and wanted it to be included in the main kernel right away. A number of framebuffer folks also expressed interest in reviewing the framebuffer integration of this driver as well.

Support for the Atmel MACB On-Chip Ethernet Module

Haavard Skinnemoen from Atmel posted a patch to support the Atmel MACB on-chip ethernet module, and listed himself as the official maintainer in the MAINTAINERS file. Jeff Garzik praised the quality of the code, saying it was a “nice, clean driver that was so painless to apply to the latest kernel”.

advertisement