

Zack's Kernel News

■ Stabilizing the core

A number of folks have been working independently to make kernel header files available to user-space programs, and to stabilize the entire kernel ABI (Application Binary Interface). A mailing list has even been set up, but it is currently lying dormant, waiting for the resumption of development after the 2.7 fork.

One result of this effort would be to allow programs compiled under one kernel version to run properly under another. This has advantages and disadvantages, since Linus Torvalds and other big developers actively discourage binary-only drivers and the like. However, it does seem like a big push will be made during the 2.7 time frame. ■

■ Live or die

With the renovation of the module-handling API in 2.5, all modules have had to be converted to use the new functions. Those that did not undergo conversion, found themselves dropped from 2.6; recently, Martin Diehl finished converting the IrDA serial dongle driver to the new module API, and David S. Miller queued it for inclusion in an upcoming 2.6 kernel release.

The renovation of the module API has been only slightly controversial, and to anyone complaining about their favorite modules disappearing from 2.6, the response has typically been, "apparently no one cares enough to port the code over!" Breakages of this kind are typically between stable series, though no one enjoys creating incompatibilities.

One side effect of changes of this kind is that it does tend to force the entire Linux software base into some kind of active maintenance; so that rarely are programs in active use, that are not also actively maintained. This is both a strength and a weakness: a strength in that the world of free software is clearly very much alive; and a weakness in that some code that is actually useful, does fall into disuse if no one steps up to maintain it. ■

■ Long support

Bjorn Ekwall has stepped down as maintainer of the drivers for the parallel port network adapters D-Link de600 and de620. These network adapters have been supported in Linux for over ten years, the de600 somewhat longer than the de620, and Bjorn has been involved with their development since the beginning; in fact he (with the help of some others) released the initial BETA de620 driver way back in June of 1994, for Linux 1.1.18.

Now, travel and busted hardware have diverted him, and the drivers are currently unsupported. Anyone interested in taking on the job should contact Bjorn directly, and he will pass the name along to David S. Miller. ■

■ Free the RAM

Jeff Dike has been petitioning for a way to free up memory used by UML throughout its session. When UML starts up, it gradually claims more and more of system RAM, as it needs to. When it doesn't need that RAM anymore, there is currently no way to communicate to the host system that the memory may be freed.

To UML it is seen as unused, but the host system can't tell the difference between RAM that UML cares about, and RAM it doesn't. So all is kept within the process. To solve this, Jeff wants to implement a way to just release RAM even when it seems like it's being used.

His proposal includes a `/dev/anon` device, which acts like anonymous memory, but actually keeps track of how many times each page is mapped, and then freeing pages whose map-count goes to zero. UML's needs in this case are fairly obscure, and there has been some resistance to the idea of `/dev/anon`.

But it's clear that either it or something similar would make a big difference to UML, so there may be similar proposals in the future, if this one doesn't pan out. ■

INFO

The Kernel Mailing List comprises the core of Linux development activities. Traffic volumes are immense 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 that take on this impossible task is Zack Brown.

Our regular monthly column keeps you up to date on the latest discussions and decisions, selected and summarized by Zack. Zack has been publishing a weekly digest, the Kernel Traffic Mailing List for several years now, reading just the digest is a time consuming task.



Linux Magazine now provides you with the quintessence of Linux Kernel activities straight from the horse's mouth.

■ Waiting for a queue

Marcel Sebek has ported the IMQ (Intermediate Queueing Device) driver from Linux 2.4 to 2.6; Patrick McHardy had maintained the 2.4 version until October of 2003, but then stepped down, asking to receive no more patches or communication about the project.

Marcel took the most recent 2.4 version, and submitted his port to David S. Miller. Patrick came out of hiding long enough to say that Marcel's work was not yet ready for inclusion, on account of the driver having bugs not yet fixed.

There's some question of whether a fully functional and fixed IMQ driver should get into the 2.6 tree, since it may be the case that other mechanisms exist to do the same job better. ■

■ Over run

A new Linux port, called Cooperative Linux, has emerged. Its purpose is to run the Linux kernel as a low-level process on a variety of other OSes, such as Microsoft Windows. Success has already been achieved under Windows 2000 and XP, and the developers, led by Dan Aloni, are optimistic that it would be possible to run on top of other OSes.

Currently, support doesn't yet exist to run Cooperative Linux under Linux itself. The closest relative to Cooperative Linux appears to be the Windows port of Jeff Dike's User-Mode Linux. ■

■ **Real Time soon**

Bernhard Kuhn has created some patches to improve real-time support in 2.6, guaranteeing latencies of no longer than 5 microseconds. Robert Love is most active in this area, and will undoubtedly have something to say about these patches before they go into the main tree.

The importance of real-time support is somewhat debated. In the past, Linus Torvalds resisted it because he felt it could only be done by creating very ugly constructs throughout the entire kernel code base. But methods were designed that satisfied him, and real-time support has progressed steadily over the years.

Gamers and music-lovers are the biggest contingent desiring real-time support; but applications that rely on real-time support include sensitive medical and scientific devices, in which any delay could literally result in fatality. ■

■ **Still life**

Willem Riede has taken on the thankless task of ide-scsi maintainership. This is a driver that Linus Torvalds has said should never have existed at all. In spite of this, various folks have found use for it, and some (including Willem) feel that the problems can all be fixed.

Some fairly significant folks, such as Bartłomiej Zolnierkiewicz, have said that ide-scsi will definitely die in the 2.7 time frame, but in spite of this, a significant number of folks want to keep it as well maintained as possible until that happens. Linus himself showed a willingness at least to consider Willem's patches, and so ide-scsi seems to be alive, at least for now. ■

■ **Moss-Covered Tortoise**

The 2.0 kernel is still being maintained. In response to a bug report from Seiichi Nakashima on the kernel mailing list, David Weinehall (the official 2.0 maintainer) was on top of the problem, and promised a new release the next week.

He made good on this, and actually released 2.0.40 on February 8, nicknaming it "The Moss-Covered Tortoise". Kernel nicknames have been common ever since the infamous "Greased Weasel", the name given by Linus Torvalds to version 1.3.51 in December 1995. Unfortunately, the Greased Weasel turned out to have some bugs, and in the process of fixing them the kernel began to slow down to the point where Linus and others began to remark on it no longer behaving much like a greased weasel at all. Since then, key kernel versions have occasionally been given names by their maintainers. ■

■ **Open for all**

Not all the mailing lists listed in the MAINTAINERS file as the proper place to submit bug reports, would allow posts by non-subscribed people. Many felt that all mailing lists receiving kernel-related bug reports should be completely open, so anyone could post.

The linux-kernel mailing list itself has always been open, with the exception of a few bannings, and some anti-spam software. A policy seems to be emerging: mailing lists listed in the MAINTAINERS file do not necessarily have to allow posts from non-members, but in that case they must provide another open mailing list and that is read by at least one main developer of that feature. ■

■ **Opening the freezer**

Marcelo Tosatti's concept of a 2.4 Deep Freeze has changed over time.

Originally, when 2.6 was about to come out, he attempted to clamp down completely, but met with stiff opposition, particularly from the XFS folks, who did eventually get their stuff into 2.4; similar stumbling blocks came up as well, which have forced Marcelo to give more warning before the actual freeze. It is looking as though bugs will delay the lock until the 2.4.27 release.

One feature not so likely to make it past the blockade is the relatively new cryptoloop feature, allowing encrypted filesystems to be mounted over loop-back. Whatever the final disposition of any particular patch, it is clear that Marcelo no longer regards 2.4 as a place for new code. As he has said, the 2.6 kernel is the proper kernel to use, to gain access to the latest features. ■

■ **In or out**

The Adaptec/DPT I2O driver failed to update itself to the new driver model, and was dropped from the 2.6 tree. Leon Toh has made some effort toward it working again. Getting the code back into the 2.6 tree will not be automatic.

Once dropped, a driver loses some of its status as 'supported', and must be considered all over again. The driver, even the 2.4 version was quite messy.

Also, it looks as though some other drivers may become available through Adaptec itself. Go Taniguchi has also been doing his own work, which has recently come under consideration. So it seems that one way or another, a workable solution is likely to be found. ■

SSV