

ZACK'S KERNEL NEWS

Idle User Alert

Alessandro Di Marco has implemented a user inactivity trigger for the 2.6.20 kernel, which issues an ACPI event when no user activity is detected for longer than a certain amount of time. Di Marco really coded this up for fun; it never occurred to him that a bunch of kernel dudes would descend upon his code with critiques and suggestions, but that's exactly what happened.

For one thing, as Arjan van de Ven pointed out, ACPI is not the right mechanism to carry the alert event and uevent would be a more logical choice. Also, Di Marco's code adds a new /proc file, which is a no-no nowadays. Pavel Machek suggested putting something in the /sys directory, perhaps in /sys/power. Machek also suggested that the whole idea would be better off in user space rather than the kernel proper, so Di Marco has been recoding his trigger to operate from user space.

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.



Kernel Summit 2007

Theodore Y. Ts'o recently initiated the planning process for the next Linux Kernel Summit, the annual invitation-only gathering of kernel hackers. Traditionally held in Ottawa, Canada, this year's gathering will be in Cambridge, England, at the request of last year's attendees. If the new location works out this time around, they may try taking the conference to additional countries in years to come.

The discussion in response to the announcement tended to focus on possible future venues. Some folks wanted Australia, some folks wanted India, and some folks wanted the Czech Republic. There were all sorts of suggestions! The travel costs for the approximately 80 invited kernel hackers seems to be a decisive factor in the choice of venue. Cheaper trips mean more employers willing to foot the bill. This in turn means there will be a tendency to hold the summit in countries that already contain or are close to as many summit attendees as possible.

So, as with most kernel projects, this one will probably go through various trials and experiments, with a rich and loudly argued history.

linux-kernel Mailing List Archives

Rob Landley noticed that the mbox-format repository of the linux-kernel mailing list didn't seem to be hosted on kernel.org anymore, and asked what had become of them. Andrew Morton said he'd be willing to upload his personal repository going back to 2000, but Dirk Behme said it would be nice to have a repository that auto-updated, so people didn't have to subscribe to the list in order to get the mail. Rob Landley also pointed out that very old lkml archives could still be found at http://www.kclug.org/old_archives/linux-activists/. Dave Jones pointed out that even earlier archives could be found at <ftp://ftp.linux.org.uk/pub/linux/history/funet-archive/> and <ftp://ftp.linux.org.uk/pub/linux/alan/Kernel/Documents/Old-Funet-Lists>.

Fixing a License Loophole

Jan Engelhardt has been working on a fix for the long-time kernel bug that allowed unscrupulous driver writers like LinuxAnt to pretend that their code was GPLed, when really it could have almost any license.

The loophole occurs when the MODULE_LICENSE is set to something like "GPL\0for files in the \"GPL\" directory; for others, only LICENSE file applies". Because of the "\0", the kernel sees "GPL" when checking for the license and then the null character fools the kernel into thinking the string has ended.

Engelhardt recently posted a patch that fixes this hole by keeping track of the length of the license text. If the length differs from the length of the license string actually read, someone may be trying to use the loophole.

There was a surprising amount of discussion around Engelhardt's patch, mainly about whether the MODULE_LICENSE feature constituted a license enforcement feature, which would violate the terms of the GNU General Public License, under which much of the kernel sources are released. That argument actually has less to do with Engelhardt's patch than with the existing MODULE_LICENSE code.

Bodo Eggert also worked on a similar patch to Engelhardt's and released it at roughly the same time. Eggert's patch also closed the \0 loophole using the same length-tracking technique as Engelhardt's patch and added several new features, such as allowing modules to have several licenses.

Eggert's patch got a lot more negative feedback than Engelhardt's – some of the additional features Eggert added were not well liked by people in the Linux community. For instance, allowing modules to have several licenses got a strong "no" vote from Alan Cox.

It may be that Bodo Eggert will be able to fix up his patch to address everyone's concerns, but for the moment it seems that Engelhardt's patch is simpler and more likely to go into the main kernel.

Syslet Asynchronous System Call Interface

Ingo Molnar has created the Syslet subsystem, a mechanism for calling system calls asynchronously from user space. Any user-space process can initiate a Syslet, and the Syslet will have access to the same run-time context as the calling routine. It actually does a bit better than being asynchronous – the user can specify whether the call should only be synchronous, only be asynchronous, or start off synchronous and then failover to be asynchronous if necessary.

Syslets are actually small user-constructed wrappers for system calls. Syslets exist entirely in kernel space, can modify user variables, and can be chained together in complex structures. Using Syslets, Molnar has measured speedups of 33.9% over cached synchronous I/O, and 19.2% over uncached synchronous I/O, which is quite a major improvement. Also, response to Syslets among

kernel developers has been very positive, to the point that a bunch of folks jumped in to help out.

Syslets look likely to go into the main-line sources eventually; among the dissenting voices, Pavel Machek points out that the Syslet wrapper mechanism amounts to an actual programming language with an in-kernel interpreter, but Linus Torvalds has made it clear that this will not be an obstacle. (Torvalds' exact words to Pavel were, "Hah. You make the joke! I get it!".)

Torvalds himself was also one of the harshest dissenters. He felt Molnar's interface was complicated and difficult for any casual user to play with. Molnar has had no real response to those criticisms, so it may be that the interface itself will have to change before Syslets make it into the official kernel.

When that will happen, or in what form that might satisfy Torvalds, is unclear.

New Intel Driver

James Ketrenos from Intel recently announced a new PRO/Wireless 3945ABG Network Connection adapter driver. Other Intel drivers have had to include a proprietary daemon.

In response to a query about this issue from Neil Brown, Ketrenos explained that Intel had made some improvements in the microcode for these devices, so that now only a microcode upgrade would be required, and no closed source daemon.

A lot of folks were impressed and pleased with this development, and with the driver itself.

It's nice to see that Intel is taking some open source issues so seriously.

Maintainership

Jon Masters has taken over module-init-tools maintainership from Rusty Russell and posted a modification to the MAINTAINERS file. Evgeniy Dushistov has added a UFS entry to the MAINTAINERS file, marked UFS as maintained, and listed himself as the maintainer.

KVM Migrating to git

Avi Kivity announced that the KVM virtual machine code would be developed under git starting immediately. He said one reason is because developers were interested in maintaining their own branches of KVM for longer periods of time before feeding their changes back to the main branch.

Advertisement