Zack's Kernel News

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.

New buckets

The controversy surrounding use of the BitKeeper proprietary version control system for kernel development continues to grow more and more heated. A number of top developers, like Alan Cox, are openly hostile to the use of BitKeeper, but this hostility has remained largely under wraps since Linus Torvalds first began using BitKeeper in early 2002. After various license changes and threats by Larry McVoy to revoke the service, it has come out in the open. It does not always express itself so violently, although there have been some notable flame wars throughout January and February.

Andrea Arcangeli finally published a script to forcibly extract data from the BitKeeper repository, not via the secret BitKeeper protocol, but via the HTTP access point Larry provides. This caused an immediate backlash from Larry, who insisted that if people used the script, he would cut off HTTP access. Andrea said that he would in fact continue to publish more such scripts, to allow non-BitKeeper users to access the repositories.

While this was going on, there were various speculations about whether BitKeeper could be legally reverse engineered in countries that had good fair-use laws, like Germany.

Finally, Pavel Machek announced the BitBucket project, intended to be a free replacement for BitKeeper. There have been a number of attempts to do something like this, notably the Subversion project, but this is the first to take the bull by the horns, and actually target the BitKeeper feature-set itself.

Surprisingly, Larry is not the only person upset by this development. A number of folks feel that any effort at a free version control system should be put into existing projects like Subversion, and not into yet another new project. On the other hand, the Subversion developers have consistently said that Subversion is not intended to replace BitKeeper, and that it has a different set of goals.

At the same time, the BitBucket developers say that they would not need to replace BitKeeper if Larry would given them equal access to BitKeeper repositories via Free tools. ■

New views

The FUSE User-Space Filesystem has reached version 1.0. A feature touted by the Hurd OS has always been the ability to create filesystem views of anything. One example has been the ability to browse an FTP site as a directory tree.

The advantage of something like this is that suddenly you can use your old familiar tools to interact with things (like FTP servers) that had always required limited, specialized applications (like FTP clients). Mounting an FTP site – or anything else – as a filesystem was the UNIX philosophy taken to its natural conclusion, and for a long time Linux was unable to do anything like this.

Things like UserFS came along, that went a long way toward giving users the ability to represent anything as a filesystem. Around the same time a lot of strange new ideas with filesystems were starting to come out. UserFS, DriverFS (now known as SysFS), DevFS and others all rose above the idea that filesystems were supposed to represent regular old files. There was a tremendous push for all these new types of filesystems to do meaningful things, like provide a replacement for the ailing ProcFS filesystem. UserFS made a few wrong assumptions, such as exposing too much of the core Virtual Filesystem subsystem API, and was eventually abandoned. Other projects along the same lines came and went, such as UVFS and IFS.

Now FUSE, by Miklos Szeredi, is trying to fill the gap left by those other attempts. It boasts ease of installation, a clean API, and the ability for regular users to do just about anything. ■

Marching forward

One of the other 'magic' filesystems, SysFS, has risen through the ranks to become one of the core interfaces to the kernel. The /proc directory, a knot of problems in the best of times, looks to be all but replaced by SysFS in the next few years. ProcFS interfaces like cpufreq are being deprecated in favor of SysFS equivalents, and new drivers are being encouraged to export their interface via SysFS exclusively, leaving /proc to fade gradually into silence. SysFS support for the ZT5550 redundant host controller is one of the most recent of these.

Other drivers exporting their own filesystems, like pcihpfs, are also migrating to SysFS, abandoning their native interfaces. The PCI hotplugging code is at the vanguard of the move toward SysFS, and has required additional SysFS features to better support hotplugging, before a true port could be adopted.

SysFS seems to be the answer to many prayers, including the proliferation of ioctls. It remains to be seen whether it can accomplish all this without sacrificing some of the clarity and cleanliness that currently sets it apart from things like ProcFS. If it fails to enforce naming and file format conventions, it could end up as a messy interface, like /proc, either locked down to an insane set of backward compatibility requirements, or else breaking them every week; waiting for something better to come along. ■

Dot dash

Speaking of insane interfaces, Tomas Szepe has continued the work on producing kernel panics in Morse code. Andrew Rodland did a bit of work on this in 2002, and Tomas brought the code up to 2.5 and cleaned it up, including new morse code table handling by Jan-Benedict Glaw. The idea is that users running the X Windowing System won't see panics on their screen.

Currently this is done by way of the on-board speaker and keyboard LEDs. So now, instead of staring helplessly at a locked screen, X users can see it on their keyboard, while enjoying the rhythms of hexadecimal notation rendered into dots and dashes on their little speaker.

Although you and I know that this feature really makes the entire system much more friendly, it may turn out to be the sort of thing we shouldn't let our parents in on quite yet. Once they've made the switch themselves, we can let them know about these features. ■

Bug Hunter

If they do hear one of these panics lilting by, and want to submit a bug report, they can read up on John Bradford's Kernel Bug Database.

In February he published a metric tonne of documentation describing his alternative to the Bugzilla database, which is currently being used for tracking kernel bugs.

John's idea is that the kernel needs a more specialized bug tracking system than Bugzilla can provide. Bugs can be searched by configuration option, which can be much more pointed than keyword searches.

This eases off the burden of categorizing each bug by hand. Since bugs are often tied to their configuration files, they are inherently categorized by the selections in that file.

The system also distinguishes between two types of bug reports, which it calls 'raw' and 'sorted'. Raw bug reports come from anyone, while sorted bug reports collect similar reports together via a volunteer contributor, eliminate those with inaccuracies, and make the whole set more useful.

Once solved, bug reports in either set are archived, and no longer appear in searches by default. It is also possible to attach patches to a given set of bug reports, so users can both contribute and test fixes.

The search for a decent bug tracking system goes back a long way, and has never been adequately solved.

Ideally, users should be able to search for exactly the symptoms they've experienced, without worrying about duplicate or ambiguous reports. They should be able to track the progress of debugging efforts, and participate if they so choose to do.

These are easy things to say, but in practice it has always proven a very tough nut to crack and hopefully this new project will help ■

advertisement