

Re: How to improve the quality of the kernel?

Re: How to improve the quality of the kernel?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-06/msg07132.html>

- *From:* "Rafael J. Wysocki" <rjw@xxxxxxx>
 - *Date:* Sun, 17 Jun 2007 20:16:16 +0200
-

On Sunday, 17 June 2007 19:42, Natalie Protasevich wrote:

On 6/17/07, Rafael J. Wysocki <rjw@xxxxxxx> wrote:

On Sunday, 17 June 2007 16:29, Adrian Bunk wrote:

On Sun, Jun 17, 2007 at 03:17:58PM +0200, Michal Piotrowski wrote:

On 17/06/07, Adrian Bunk
<bunk@xxxxxxxxxx> wrote:

...

Fine with me, but:

There are not so simple cases like big infrastructure patches with 20 other patches in the tree depending on it causing a regression, or even worse, a big infrastructure patch exposing a latent old bug in some completely different area of the kernel.

It is different case.

"If the patch introduces a new regression"

introduces != exposes an old bug

My remark was meant as a note "this sentence can't handle all regressions" (and for a user it doesn't matter whether a new

Re: How to improve the quality of the kernel?

regression is introduced or an old regression exposed).

It could be we simply agree on this one. ;-)

Removal of 20 patches will be painful, but sometimes you need to "choose minor evil to prevent a greater one" [1].

And we should be aware that reverting is only a workaround for the real problem which lies in our bug handling.

...

And this is something I want to emphasize again.

How can we make any progress with the real problem and not only the symptoms?

I think that we can handle bug reports like we handle modifications of code.

Namely, for each subsystem there can be a person (or a team) responsible for handling bugs, by which I don't mean fixing them, but directing bug reports at the right developers or subsystem maintainers, following the history of each bug report etc. [Of course, these people can choose to use the bugzilla or any other bug tracking system they want, as long as it works for them.]

The email addresses of these people should be known (and even documented), so that everyone can notify them if need be and so that it's clear who should handle given bug reports.

Just an idea. :-)

Those are very good ideas indeed. The whole development process came to the point when all realize that something needs to be done for the team to balance out new development and old and recent unresolved issues that are piling up...

Re: How to improve the quality of the kernel?

I've looked through a number of bugzillas recently and here is my scoop on shortcomings and some ideas. I am not sure how realistic they are, probably might fall into "wishful thinking" category.

The way bugs get tracked and resolved is definitely a "no guarantee", and main reasons are:
not enough time for a maintainer to attend to them all
nobody else (except at best very few busy people) knows about majority of the problems. Andrew and Adrian and Michal post the most pressing ones. But there are many many smaller ones that are not assessed and not being taken care of.
many problems are not easily reproducible and not easy to verify because there is no identical system, motherboard, application, etc.
in case if reporter doesn't stick around until the end of the bug's life.

I agree. In addition, there is only a limited time window in which it makes sense to debug given problem before the kernel changes too much (that of course depends on the subsystem in question).

Maybe along with bugzilla there should be another tracking tool – for resources and systems that are available to individual developers.

Yes, that would be very nice to have.

Someone might have same or similar system to verify fixes in case if the reporter disappears or "the system is gone now". Requests for specific hardware can be automatically generated by the bugzilla say. Those can be posted once in a while for everyone to see and chip in and acknowledge if they happen to have such hardware and able to run a quick test to at least verify the patch. Statistically, such need doesn't happen often for each type of hardware, so it shouldn't be a big burden for owners.

Besides, the database and resources can be useful for developers who want to test their new patches on variety of hardware. This might prevent future regressions which often caused by lack of testing as we all know.

For that, I think, some "professional testers" would be needed ...

Greetings,
Rafael

—

Re: How to improve the quality of the kernel?

"Premature optimization is the root of all evil." – Donald Knuth

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>