



Finalmente disponibile nuovo kernel 2.6.30 rilascio da <http://www.kernel.org/>

Ricordo che questo kernel è usabile con attenzione ma che manca di tutte le patch e moduli 3rdparty tipiche dei kernel Mandriva non essendo quello ufficiale di Mandriva.

commit 96050dfb25966612008dcea7d342e91fa01e993c

char: mxser, fix ISA board lookup

There's a bug in the mxser kernel module that still appears in the 2.6.29.4 kernel.

mxser_get_ISA_conf takes a ioaddress as its first argument, by passing the not of the ioaddr, you're effectively passing 0 which means it won't be able to talk to an ISA card. I have tested this, and removing the ! fixes the problem.

commit a61d90d75d0f9e86432c45b496b4b0fbf0fd03dc

jbd: fix race in buffer processing in commit code

In commit code, we scan buffers attached to a transaction. During this scan, we sometimes have to drop `j_list_lock` and then we recheck whether the journal buffer head didn't get freed by `journal_try_to_free_buffers()`.

But checking for `buffer_jbd(bh)` isn't enough because a new journal head could get attached to our buffer head. So add a check whether the journal head remained the same and whether it's still at the same transaction and list.

This is a nasty bug and can cause problems like memory corruption (use after free) or trigger various assertions in JBD code (observed).

commit 463aea1a1c49f1a7d4b50656dfd6c8bb33358b1b

autofs4: remove hashed check in `validate_wait()`

The recent `->lookup()` deadlock correction required the directory inode mutex to be dropped while waiting for expire completion. We were concerned about side effects from this change and one has been identified.

I saw several error messages.

They cause autofs to become quite confused and don't really point to the actual problem.

Things like:

```
handle_packet_missing_direct:1376: can't find map entry for (43,1827932)
```

which is usually totally fatal (although in this case it wouldn't be except that I treat it as such because it normally is).

```
do_mount_direct: direct trigger not valid or already mounted
/test/nested/g3c/s1/ss1
```

which is recoverable, however if this problem is at play it can cause autofs to become quite confused as to the dependencies in the mount tree because mount triggers end up mounted multiple times. It's hard to accurately check for this over mounting case and automount shouldn't need to if the kernel module is doing its job.

There was one other message, similar in consequence of this last one but I can't locate a log example just now.

When checking if a mount has already completed prior to adding a new mount request to the wait queue we check if the dentry is hashed and, if so, if it is a mount point. But, if a mount successfully completed while we

slept on the wait queue mutex the dentry must exist for the mount to have completed so the test is not really needed.

Mounts can also be done on top of a global root dentry, so for the above case, where a mount request completes and the wait queue entry has already been removed, the hashed test returning false can cause an incorrect callback to the daemon. Also, `d_mountpoint()` is not sufficient to check if a mount has completed for the multi-mount case when we don't have a real mount at the base of the tree.

Potete scaricarlo da [qui](#), al solito tasto destro salva destinazione con nome.

Ecco il [log](#) dei fix.

ciao.

grissino.