



Disponibile versione stabile nuovo kernel 2.6.30.1 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.

bsdacct: fix access to invalid filp in acct\_on()

commit df279ca8966c3de83105428e3391ab17690802a9 upstream.

The file opened in acct\_on and freshly stored in the ns->bacct struct can be closed in acct\_file\_reopen by a concurrent call after we release acct\_lock and before we call mntput(file->f\_path.mnt).

Record file->f\_path.mnt in a local variable and use this variable only.

commit 6e9b0f6d101c80eb88a648d6746031ca77043043

Author: Felix Blyakher <felixb@sgi.com>

Date: Thu Jun 11 17:07:28 2009 -0500

xfs: fix freeing memory in xfs\_getbmap()

commit 7747a0b0af5976ba3828796b4f7a7adc3bb76dbd upstream.

Regression from commit 28e211700a81b0a934b6c7a4b8e7dda843634d2f.  
Need to free temporary buffer allocated in xfs\_getbmap().

commit e1b41bb33a93eb3071c27db2c7313d72c5f2500c

KVM: x86: silence preempt warning on kvm\_write\_guest\_time

commit 2dea4c84bc936731668b5a7a9fba5b436a422668 upstream.

This issue just appeared in kvm-84 when running on 2.6.28.7 (x86-64)  
with PREEMPT enabled.

We're getting syslog warnings like this many (but not all) times qemu  
tells KVM to run the VCPU:

BUG: using smp\_processor\_id() in preemptible [00000000] code:

qemu-system-x86/28938

caller is kvm\_arch\_vcpu\_ioctl\_run+0x5d1/0xc70 [kvm]

Pid: 28938, comm: qemu-system-x86 2.6.28.7-mtyrel-64bit

Call Trace:

debug\_smp\_processor\_id+0xf7/0x100

kvm\_arch\_vcpu\_ioctl\_run+0x5d1/0xc70 [kvm]

? \_\_wake\_up+0x4e/0x70

? wake\_futex+0x27/0x40

kvm\_vcpu\_ioctl+0x2e9/0x5a0 [kvm]

enqueue\_hrtimer+0x8a/0x110

\_spin\_unlock\_irqrestore+0x27/0x50

vfs\_ioctl+0x31/0xa0

do\_vfs\_ioctl+0x74/0x480

sys\_futex+0xb4/0x140

sys\_ioctl+0x99/0xa0

system\_call\_fastpath+0x16/0x1b

As it turns out, the call trace is messed up due to gcc's inlining, but  
I isolated the problem anyway: kvm\_write\_guest\_time() is being used in a  
non-thread-safe manner on preemptable kernels.

Basically kvm\_write\_guest\_time()'s body needs to be surrounded by  
preempt\_disable() and preempt\_enable(), since the kernel won't let us  
query any per-CPU data (indirectly using smp\_processor\_id()) without  
preemption disabled. The attached patch fixes this issue by disabling  
preemption inside kvm\_write\_guest\_time().

[marcelo: surround only \_\_get\_cpu\_var calls since the warning

is harmless]

commit 06ea0c9a0bb0d882a436ecf7876a08ace86de9a8

drm/i915: correct suspend/resume ordering

commit 9e06dd39f2b6d7e35981e0d7aded618686b32ccb upstream.

We need to save register state *after* idling GEM, clearing the ring, and uninstalling the IRQ handler, or we might end up saving bogus fence regs, for one. Our restore ordering should already be correct, since we do GEM, ring and IRQ init after restoring the last register state, which prevents us from clobbering things.

I put this together to potentially address a bug, but I haven't heard back if it fixes it yet. However I think it stands on its own, so I'm sending it in.

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

Ecco il [log](#) dei fix.

ciao.

grissino.