

Paul Menage wrote:

> On 11/23/06, Pavel Emelianov <xemul@openvz.org> wrote:

>> You mean moving is like this:

>>

>> old\_bc = task->real\_bc;

>> task->real\_bc = new\_bc;

>> cmpxchg(&tsk->exec\_bc, old\_bc, new\_bc);

>>

>> ? Then this won't work:

>>

>> Initialisation:

>> current->exec\_bc = init\_bc;

>> current->real\_bc = init\_bc;

>> ...

>> IRQ:

>> current->exec\_bc = init\_bc;

>> ...

>> old\_bc = tsk->real\_bc; /\* init\_bc \*/

>> tsk->real\_bc = bc1;

>> cx(tsk->exec\_bc, init\_bc, bc1); /\* ok \*/

>> ...

>> Here at the middle of an interrupt

>> we have bc1 set as exec\_bc on task

>> which IS wrong!

>

> You could get round that by having a separate "irq\_bc" that's never

> valid for a task not in an interrupt.

No no no. This is not what is needed. You see, we do have to set exec\_bc as temporary (and atomic) context. Having temporary context is 1. flexible 2. needed by beancounters' network accountig. We have to track this particular scenario.

Moreover making get\_exec\_bc() as

if (in\_interrupt())

return &irq\_bc;

else

return current->exec\_bc;

is awful. It must me simple and stupid to allow us making temporary contexts in any place of code.

Maybe we can make smth similar to wait\_task\_inactive and change it's beancounter before unlocking the runqueue?

> Paul

>

---