
Subject: Re: [ckrm-tech] [PATCH 4/13] BC: context handling
Posted by [Paul Menage](#) on Thu, 23 Nov 2006 08:53:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

```
>
> We can do the following:
>
> if (tsk == current)
>     /* fast way */
>     tsk->exec_bc = bc;
> else
>     /* slow way */
>     stop_machine_run(...);
>
> What do you think?
```

How about having two pointers per task:

- exec_bc, which is the one used for charging
- real_bc, which is the task's actual beancounter

at the start of irq, do

```
current->exec_bc = &init_bc;
```

at the end of irq, do

```
current->exec_bc = current->real_bc;
```

When moving a task to a different bc do:

```
task->real_bc = new_bc;
atomic_cmpxchg(&task->exec_bc, old_bc, new_bc);
```

(with appropriate memory barriers). So if the task is in an irq with a modified exec_bc pointer, we do nothing, otherwise we update exec_bc to point to the new real_bc.

Paul
