

---

Subject: Re: [ckrm-tech] [PATCH 5/13] BC: user interface (syscalls)

Posted by [Balbir Singh](#) on Wed, 06 Sep 2006 08:57:38 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Pavel Emelianov wrote:

> Balbir Singh wrote:

>>> +

>>> +asmlinkage long sys\_set\_bcid(bcid\_t id)

>>> +{

>>> + int error;

>>> + struct beancounter \*bc;

>>> + struct task\_beancounter \*task\_bc;

>>> +

>>> + task\_bc = &current->task\_bc;

>> I was playing around with the bc patches and found that to make

>> use of bc's, I had to actually call set\_bcid() and then exec() a

>> task/shell so that the id would stick around. Would you consider

> That sounds very strange as sys\_set\_bcid() actually changes current's

> exec\_bc.

> One note is about mm's bc - mm obtains new bc only after fork or exec -

> that's

> true. But kmemsize starts charging right after the sys\_set\_bcid.

I was playing around only with kmemsize. I think the reason for my observation is this

bash --> (my utility) --> set\_bcid()

Since bash spawns my utility in a separate process, it creates and assigns a bean counter to it and then my utility exits. Unless it spawns/exec()'s a new shell, the beancounter is freed when the task exits (my utility).

>> changing sys\_set\_bcid to sys\_set\_task\_bcid() or adding a new

>> system call sys\_set\_task\_bcid()? We could pass the pid that we

>> intend to associate with the new id. This also means we'll need

>> locking around to protect task->task\_bc.

>

--

Balbir Singh,  
Linux Technology Center,  
IBM Software Labs

---