

---

Subject: Re: Re: [RFC][PATCH 7/7] UBC: proc interface

Posted by [kir](#) on Thu, 17 Aug 2006 16:12:30 GMT

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

---

Greg KH wrote:

> On Thu, Aug 17, 2006 at 05:43:16PM +0400, Kirill Korotaev wrote:

>

>>> On Wed, Aug 16, 2006 at 07:44:30PM +0400, Kirill Korotaev wrote:

>>>

>>>

>>>> Add proc interface (/proc/user\_beancounters) allowing to see current  
>>>> state (usage/limits/fails for each UB). Implemented via seq files.

>>>>

>>> Ugh, why /proc? This doesn't have anything to do with processes, just  
>>> users, right? What's wrong with /sys/kernel/ instead?

>>>

>> We can move it, if there are much objections.

>>

>

> I am objecting. /proc is for processes so do not add any new files  
> there that do not deal with processes.

>

>

>>> Or /sys/kernel/debug/user\_beancounters/ in debugfs as this is just a  
>>> debugging thing, right?

>>>

>> debugfs is usually OFF imho.

>>

>

> No, distros enable it.

>

>

>> you don't export meminfo information in debugfs, correct?

>>

>

> That is because the meminfo is tied to processes, or was added to proc  
> before debugfs came about.

>

> Then how about just /sys/kernel/ instead and use sysfs? Just remember,  
> one value per file please.

>

I see two problems with that. But let me first describe the current  
/proc/user\_beancounters. This is how it looks like from inside a container:

```
# cat /proc/user_beancounters
```

```
Version: 2.5
```

uid	resource	held	maxheld	barrier	limit	failcnt
123:	kmemsize	836919	1005343	2752512	2936012	0

```
lockedpages      0      0      32      32      0
privvmpages     4587    7289    49152    53575    0
.....(more lines like that).....
```

I.e. a container owner can take a glance over the current parameters, their usage and (the thing that is really important) fail counters. Fail counter increases each time a parameter hits the limit. This is very straightforward way for container's owner to see if everything is OK or not.

So, the problems with /sys are:

(1) Getting such info from 40+ files requires at least some script, while now cat is just fine.

(2) Do we want to virtualize sysfs and enable /sys for every container? Note that user\_beancounters statistics is really needed for container's owner to see. At the same time, container's owner should not be able to modify it -- so we should end up with read/write ubc entries for the host system and read-only ones for the container.

Taking into account those two issues, current /proc/user\_beancounters might be not that bad.

---