
Subject: Re: How to draw values for /proc/stat
Posted by [KOSAKI Motohiro](#) on Sun, 11 Dec 2011 19:11:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

>>> IOW a /proc namespace coupled to cgroup scope would do what you want.
>>> Now my head hurts..
>>
>> Mine too. The idea is good, but too broad. Boils down to: How do you
>> couple them? And none of the methods I thought about seemed to make any
>> sense.
>>
>> If we really want to have the values in /proc being opted-in, I think
>> Kamezawa's idea of a mount option is the winner so far.

```
> diff --git a/include/linux/cgroup.h b/include/linux/cgroup.h
> index 1b7f9d5..f0bc2e9 100644
> --- a/include/linux/cgroup.h
> +++ b/include/linux/cgroup.h
> @@ -158,6 +158,7 @@ enum {
>  * Clone cgroup values when creating a new child cgroup
>  */
> CGRP_CLONE_CHILDREN,
> + CGRP_PROC_OVERLAY,
> };
```

I'm not cgroup expert, but I doubt it is mount option. I suspect it's cgroup option. That's said, if we have following two directories,

/cgroup-for-virtualization
/cgroup-for-resource-management

are both directory affected the overlay flag? I don't think it is not optimal. Why? we must care some system software (e.g. kvm, systemd) are using cgroup internally and we expected this trend will grow more.

So, I doubt namespace issue can be solved by such tiny patch.
