Subject: Re: [ckrm-tech] [RFC] Resource Management - Infrastructure choices
Posted by Pavel Emelianov on Wed, 01 Nov 2006 08:01:31 GMT
View Forum Message <> Reply to Message

[snip]

>> 2. Having configfs as the only interface doesn't alow
>>   people having resource controll facility w/o configfs.
>>   Resource controller must not depend on any "feature".
>
> One flexibility configfs (and any fs-based interface) offers is, as Matt
> had pointed out sometime back, the ability to delage management of a
> sub-tree to a particular user (without requiring root permission).
>
> For ex:
>
>   /
> |
>   -----------------
> | |
>   vatsa (70%) linux (20%)
> |
>   ---------------------------------
> | | |
>   browser (10%)   compile (50%)   editor (10%)
>
> In this, group 'vatsa' has been alloted 70% share of cpu. Also user
> 'vatsa' has been given permissions to manage this share as he wants. If
> the cpu controller supports hierarchy, user 'vatsa' can create further
> sub-groups (browser, compile ..etc) -without- requiring root access.

I can do the same using bcctl tool and sudo :)

> Also it is convenient to manipulate resource hierarchy/parameters thr a
> shell-script if it is fs-based.
>
>> 3. Configfs may be easily implemented later as an additional
>>   interface. I propose the following solution:
>
> Ideally we should have one interface - either syscall or configfs - and
> not both.

Agree.

> Assuming your requirement of auto-deleting objects in configfs can be
> met thr' something similar to cpuset's notify_on_release, what other
> killer problem do you think configfs will pose?
>

>
>>> - Should we have different groupings for different resources?
>> This breaks the idea of groups isolation.
>
> Sorry dont get you here. Are you saying we should support different
> grouping for different controllers?

Not me, but other people in this thread.
