Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)
Posted by Balbir Singh on Wed, 06 Sep 2006 19:17:56 GMT
View Forum Message <> Reply to Message

Kirill Korotaev wrote:
> Balbir Singh wrote:
>> Kirill Korotaev wrote:
>>
>>> Core Resource Beancounters (BC) + kernel/user memory control.
>>>
>>> BC allows to account and control consumption
>>> of kernel resources used by group of processes.
>>>
>>> Draft UBC description on OpenVZ wiki can be found at
>>> http://wiki.openvz.org/UBC_parameters
>>>
>>> The full BC patch set allows to control:
>>> - kernel memory. All the kernel objects allocatable
>>> on user demand should be accounted and limited
>>> for DoS protection.
>>> E.g. page tables, task structs, vmas etc.
>>
>> One of the key requirements of resource management for us is to be able to
>> migrate tasks across resource groups. Since bean counters do not associate
>> a list of tasks associated with them, I do not see how this can be done
>> with the existing bean counters.
> It was discussed multiple times already.
> The key problem here is the objects which do not _belong_ to tasks.
> e.g. IPC objects. They exist in global namespace and can't be reaccounted.
> At least no one proposed the policy to reaccount.
> And please note, IPCs are not the only such objects.
>
> But I guess your comment mostly concerns user pages, yeah?

Yes.

> In this case reaccounting can be easily done using page beancounters
> which are introduced in this patch set.
> So if it is a requirement, then lets cooperate and create such functionality.
>

Sure, let's cooperate and talk.

> So for now I see 2 main requirements from people:
> - memory reclamation
> - tasks moving across beancounters
>

Some not quite so urgent ones - like support for guarantees. I think this can be worked out as we make progress.

> I agree with these requirements and lets move into this direction.
> But moving so far can't be done without accepting:
> 1. core functionality
> 2. accounting
>

Some of the core functionality might be a limiting factor for the requirements. Lets agree on the requirements, I think its a great step forward and then build the core functionality with these requirements in mind.

> Thanks,
> Kirill
>

--

Balbir Singh,
Linux Technology Center,
IBM Software Labs