Subject: Re: containers development plans (July 10 version)
Posted by Balbir Singh on Wed, 11 Jul 2007 06:31:57 GMT
View Forum Message <> Reply to Message

Paul Menage wrote:
> On 7/10/07, Serge E. Hallyn <serge@hallyn.com> wrote:
>>
>> A (still under construction) list of features we expect to be worked on
>> next year looks like this:
>> 4. task containers functionality
>> specific containers
>
> A couple of more container subsystem requests that have come out of
> the Linux Foundation Japan symposium, although I think they've also
> been mentioned before more than once - per-container swap and disk I/O
> scheduling.
>

I think per container swap is interesting

> I'm not familiar enough with the current Linux disk scheduler code to
> know how easy/hard it is to add rate guarantees on a per-container
> basis, but the swap one should be easier.
>
> One potential issue with the swap container is how integrated should
> it be with the memory controller? I can certainly see people wanting
> to be able to use a swap controller without requiring a page-based
> memory controller (e.g. you might want to combine it with node-based
> control via cpusets instead) but adding two pointers to the mm_struct,
> one for swap controller subsystem and one for memory controller
> subsystem, seems a little bit ugly.
>

Well, it depends on how you define ugly. We could so something like
the namespace approach, have something like

struct mem_container_ptrs {
 swap_list;
 mem_container_ptr;
};

Although, I agree that per container swap is important, I feel that
we should add in the functionality, once we have basic page based
memory controller. It would make the whole setup easier to test
for functionality and performance.

> Paul

--
Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL