Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by Andrea Righi on Thu, 26 Jul 2012 11:16:44 GMT
View Forum Message <> Reply to Message

On Wed, Jul 25, 2012 at 02:00:41PM +0400, Glauber Costa wrote:
> On 07/25/2012 02:00 PM, Eric W. Biederman wrote:
> > Glauber Costa <glommer@parallels.com> writes:
> >
> >> On 07/12/2012 01:41 AM, Kir Kolyshkin wrote:
> >>> Gentlemen,
> >>>
> >>> We are organizing containers mini-summit during next Linux Plumbers (San
> >>> Diego, August 29-31).
> >>> The idea is to gather and discuss everything relevant to namespaces,
> >>> cgroups, resource management,
> >>> checkpoint-restore and so on.
> >>>
> >>> We are trying to come up with a list of topics to discuss, so please
> >>> reply with topic suggestions, and
> >>> let me know if you are going to come.
> >>>
> >>> I probably forgot a few more people (such as, I am not sure who else
> >>> from Google is working
> >>> on cgroups stuff), so fill free to forward this to anyone you believe
> >>> should go,
> >>> or just let me know whom I missed.
> >>>
> >>> Regards,
> >>> Kir.
> >>
> >> BTW, sorry for not replying before (vacations + post-vacations laziness)
> >>
> >> I would be interested in adding /proc virtualization to the discussion.
> >> By now it seems userspace would be the best place for that to happen, in
> >> a fuse overlay. I know Daniel has an initial implementation of that, and
> >> it would be good to have it as library that both OpenVZ and LXC (and
> >> whoever else wants) can use.
> >>
> >> Shouldn't take much time...
> >
> > What would you need proc virtualization for?
> >
>
> proc provides a lot of information that userspace tools rely upon.
> For instance, when running top, you can draw per-process figures from
> what we have now, but you can't make sense of percentages without
> aggregating container-wide information.
>

> When you read /proc/cpuinfo, as well, you would expect to see something
> that matches your container configuration.
>
> "free" is another example. The list go on.

Another interesting feature IMHO would be the per-cgroup loadavg. A typical use case could be a monitoring system that wants to know which containers are more overloaded than others, instead of using a single system-wide measure in /proc/loadavg.

-Andrea
