
Subject: Re: OLS paper topics
Posted by [serue](#) on Wed, 30 Jan 2008 15:52:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Balbir Singh (balbir@linux.vnet.ibm.com):

> Cedric Le Goater wrote:
> >
> >>> == Topic ==
> >>>
> >>> Namespaces, containers, cgroups, and container checkpoint/restart.
> >>>
> >>> == Description ==
> >>>
> >>> Development for namespaces and cgroups is well underway in the
> >>> mainline kernel. To keep momentum going and keep the loosely
> >>> knit teams of developers well-coordinated, a physical meeting in
> >>> which to discuss future development plans is needed.
> >>> Additionally, we are at a point where crucial decisions about
> >>> the nature of a "container object" and about the checkpoint/restart
> >>> design need to be made.
> >>>
> >>> A final set of topics will be decided upon through mailing lists
> >>> ahead of time, but potential topics include:
> >>> * Handling filesystem/namespace synchronization
> >>> * Handling of /proc and /sysfs within containers
> >>> * Additional needed namespaces (i.e. device namespace)
> >>> * Nature of a 'container' ??? kernel object or userspace fiction
> >>> * Additional cgroups and their design
> >>> * How to initiate and synchronize checkpoint/restart
> >>> * How to enter a 'container' ?
> >>>
> >>
> >> Resource Management for containers?
> >
> > Yes we need that. There are a few possible conflicts in requirements
> > that need to be discussed. For example, a resource management req
> > would be to be able to move a process from one group to another and C/R
> > wouldn't.
>
> Cedric,
>
> Why is that a conflict w.r.t resource management, isn't that more of a cgroup
> feature - task migration.

Oh, no - I think Cedric was saying that tasks need to be moved from container2 to container3 on the system for resourcemgmt - that is, reclassified, not migrated.

But for task migration we have resources tied to containers and processes are using those resources, so we can't migrate those tasks to another container in that case.

That's probably solved at the cgroup->container bindings, no? i.e. ns_cgroup 2854 is my container and it's in a cpuset_cgroup tied to cpus 2 and 3, now i want to open it up to cpus 1-4, so i leave the tasks in ns_cgroup 2854 but move it to another cpuset...

> We need task migration for resource management, but
> task migration is a system administrator/management function. I don't see the
> conflict, but may be I am missing something very obvious.

>

> >>> == Moderators ==

> >>> {{Out|These are just suggestions, please edit here or let us know if you
> >>> do or do not want to moderate.}}

> >>>

> >>> * Namespaces/containers: Serge Hallyn, Cedric Le Goater

> >>> * cgroups: Paul Menage

> >>> * Checkpoint/restart: Dave Hansen, Cedric Le Goater

> >

> > IMO, 2 is a minimum to get the work done (it will require a minimum of
> > organization) but the moderators don't necessarily have to be from the
> > same team.

> >

> > Or, having some one in charge of the logistics would be better ? In that
> > case, you can remove me from the moderators.

> >

> >> Resource Management? - I don't mind moderating that if it happens to be
> >> included in the summit or can we merge it with cgroups?

> >

> > Good thing.

> >

> > Cheers,

> >

> > C.

>

>

> --

> Warm Regards,

> Balbir Singh

> Linux Technology Center

> IBM, ISTL

>

>

>

> _____
> Containers mailing list

> Containers@lists.linux-foundation.org

> <https://lists.linux-foundation.org/mailman/listinfo/containers>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
