Subject: Re: OLS paper topics
Posted by Balbir Singh on Wed, 30 Jan 2008 16:54:59 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> 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.
>

OK, so the concern is with automatic reclassification

> 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...
>

Yes, I agree

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

_____