

On Wed, 2006-09-20 at 09:25 -0700, Christoph Lameter wrote:

> On Tue, 19 Sep 2006, Rohit Seth wrote:

>

> > For example, a user can run a batch job like backup inside containers.

> > This job if run unconstrained could step over most of the memory present

> > in system thus impacting other workloads running on the system at that

> > time. But when the same job is run inside containers then the backup

> > job is run within container limits.

>

> I just saw this for the first time since linux-mm was not cced. We have

> discussed a similar mechanism on linux-mm.

>

> We already have such a functionality in the kernel its called a cpuset. A

Christoph,

There had been multiple discussions in the past (as recent as Aug 18, 2006), where we (Paul and CKRM/RG folks) have concluded that cpuset and resource management are orthogonal features.

cpuset provides "resource isolation", and what we, the resource management guys want is work-conserving resource control.

cpuset partitions resource and hence the resource that are assigned to a node is not available for other cpuset, which is not good for "resource management".

chandra

PS:

Aug 18 link: <http://marc.theaimsgroup.com/?l=linux-kernel&m=115593114408336&w=2>

Feb 2005 thread: <http://marc.theaimsgroup.com/?l=ckrm-tech&m=110790400330617&w=2>

> container could be created simply by creating a fake node that then

> allows constraining applications to this node. We already track the

> types of pages per node. The statistics you want are already existing.

> See /proc/zoneinfo and /sys/devices/system/node/node*/*.

>

> > We use the term container to indicate a structure against which we track

> > and charge utilization of system resources like memory, tasks etc for a

> > workload. Containers will allow system admins to customize the

> > underlying platform for different applications based on their

> > performance and HW resource utilization needs. Containers contain

> > enough infrastructure to allow optimal resource utilization without
> > bogging down rest of the kernel. A system admin should be able to
> > create, manage and free containers easily.
>
> Right that's what cpusets do and it has been working fine for years. Maybe
> Paul can help you if you find anything missing in the existing means to
> control resources.
>
> -----
> Take Surveys. Earn Cash. Influence the Future of IT
> Join SourceForge.net's Techsay panel and you'll get the chance to share your
> opinions on IT & business topics through brief surveys -- and earn cash
> [http://www.techsay.com/default.php?page=join.php&p=sourc eforge&CID=DEVDEV](http://www.techsay.com/default.php?page=join.php&p=sourc%20eforge&CID=DEVDEV)
> -----
> ckrm-tech mailing list
> <https://lists.sourceforge.net/lists/listinfo/ckrm-tech>
--

Chandra Seetharaman | Be careful what you choose....
- sekharan@us.ibm.com |you may get it.
