
Subject: Re: [patch00/05]: Containers(V2)- Introduction
Posted by [Nick Piggin](#) on Wed, 20 Sep 2006 17:07:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Sep 20, 2006 at 09:48:13AM -0700, Christoph Lameter wrote:

> On Wed, 20 Sep 2006, Nick Piggin wrote:

>

> > > Right thats 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.

> >

> > What I like about Rohit's patches is the page tracking stuff which

> > seems quite simple but capable.

> >

> > I suspect cpusets don't quite provide enough features for non-exclusive

> > use of memory (eg. page tracking for directed reclaim).

>

> Look at the VM statistics please. We have detailed page statistics per

> zone these days. If there is anything missing then this would best be put

> into general functionality. When I looked at it, I saw page statistics

> that were replicating things that we already track per zone. All these

> would become available if a container is realized via a node and we would

> be using proven VM code.

Look at what the patches do. These are not only for hard partitioning of memory per container but also those that share memory (eg. you might want each to share 100MB of memory, up to a max of 80MB for an individual container).

The nodes+cpusets stuff doesn't seem to help with that because you with that because you fundamentally need to track pages on a per container basis otherwise you don't know who's got what.

Now if, in practice, it turns out that nobody really needed these features then of course I would prefer the cpuset+nodes approach. My point is that I am not in a position to know who wants what, so I hope people will come out and discuss some of these issues.
