
Subject: Re: [patch00/05]: Containers(V2)- Introduction
Posted by [Nick Piggin](#) on Wed, 20 Sep 2006 17:30:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

Alan Cox wrote:

> Ar lau, 2006-09-21 am 03:00 +1000, ysgrifennodd Nick Piggin:
>
>> > I've been thinking a bit on that problem, and it would be possible to
>> > share all address_space pages equally between attached containers, this
>> > would lose some accuracy, since one container could read 10% of the file
>> > and another 90%, but I don't think that is a common scenario.
>>
>>
>> Yeah, I'm not sure about that. I don't think really complex schemes
>> are needed... but again I might need more knowledge of their workloads
>> and problems.
>
>
> Any scenario which permits "cheating" will be a scenario that happens
> because people will try and cheat.

That's true, and that's one reason why I've advocated the solution implemented by Rohit's patches, that is: just throw in the towel and be happy to count just pages.

Look at the beancounter stuff, and it has hooks (in the form of gfp flags) throughout the tree, and they still manage to miss accounting user exploitable memory overallocation from some callers. Maintaining that will be much more difficult and error prone.

--

SUSE Labs, Novell Inc.

Send instant messages to your online friends <http://au.messenger.yahoo.com>
