

---

Subject: Re: Linux-VServer example results for sharing vs. separate mappings ...

Posted by [Bodo Eggert](#) on Thu, 29 Mar 2007 07:24:43 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Ethan Solomita <solo@google.com> wrote:

> I suggest a variant on what Andrew says: don't change reclaim.  
> Instead, when referencing a page, don't mark the page as referenced if  
> the current task is not permitted to allocate from the page's node. I'm  
> thinking in terms of cpusets, with each task having a nodemask of  
> mems\_allowed. This may result in a page being thrown out unnecessarily  
> and brought back in from disk, but when memory is tight that is what  
> happens. An optimization might be to keep track of who is referencing  
> the page and migrate it to their memory instead of reclaiming it, but  
> that would require reclaim to know the task/cpuset/container of the  
> referencing task.

If you are near reclaim and the page from that other node hasn't been touched recently, maybe you could steal it. Off cause this is only applicable if the checks are cheap enough.

--

Funny quotes:

40. Isn't making a smoking section in a restaurant like making a peeing section in a swimming pool?

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---