

---

Subject: Re: [RFC][PATCH 2/7] RSS controller core  
Posted by [Dave Hansen](#) on Fri, 16 Mar 2007 19:46:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Fri, 2007-03-16 at 12:54 -0600, Eric W. Biederman wrote:  
> Dave Hansen <[hansendc@us.ibm.com](mailto:hansendc@us.ibm.com)> writes:  
> > <http://linux-mm.org/SoftwareZones>

> Looking at your page, and I'm too lazy to figure out how to update it  
> I have a couple of comments.

You just need to create an account by clicking the Login button. It lets you edit things after that. But, I'd be happy to put anything in there you see fit.

> - Why do limits have to apply to the unmapped page cache?

To me, it is just because it consumes memory. Unmapped cache is, of course, much more easily reclaimed than mapped files, but it still fundamentally causes pressure on the VM.

To me, a process sitting there doing constant reads of 10 pages has the same overhead to the VM as a process sitting there with a 10 page file mapped, and reading that.

> - Could you mention proper multi process RSS limits.  
> (I.e. we count the number of pages each group of processes have mapped  
> and limit that).  
> It is the same basic idea as partial page ownership, but instead of  
> page ownership you just count how many pages each group is using and  
> strictly limit that. There is no page ownership or partial charges.  
> The overhead is just walking the rmap list at map and unmap time to  
> see if this is the first users in the container. No additional kernel  
> data structures are needed.

I've tried to capture this. Let me know what else you think it needs.

<http://linux-mm.org/SoftwareZones>

-- Dave

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---