

---

Subject: Re: [RFC][PATCH 4/7] RSS accounting hooks over the code  
Posted by [Vaidyanathan Srinivas](#) on Wed, 14 Mar 2007 14:43:12 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Nick Piggin wrote:

> Vaidyanathan Srinivasan wrote:

>  
>> Accounting becomes easy if we have a container pointer in struct page.  
>> This can form base ground for building controllers since any memory  
>> related controller would be interested in tracking pages. However we  
>> still want to evaluate if we can build them without bloating the  
>> struct page. Pagecache controller (2) we can implement with container  
>> pointer in struct page or container pointer in struct address space.  
>  
> The thing is, you have to worry about actually getting anything in the  
> kernel rather than trying to do fancy stuff.  
>  
> The approaches I have seen that don't have a struct page pointer, do  
> intrusive things like try to put hooks everywhere throughout the kernel  
> where a userspace task can cause an allocation (and of course end up  
> missing many, so they aren't secure anyway)... and basically just  
> nasty stuff that will never get merged.  
>  
> Struct page overhead really isn't bad. Sure, nobody who doesn't use  
> containers will want to turn it on, but unless you're using a big PAE  
> system you're actually unlikely to notice.  
>  
> But again, I'll say the node-container approach of course does avoid  
> this nicely (because we already can get the node from the page). So  
> definitely that approach needs to be discredited before going with this  
> one.

I agree :)

>> Building on this patchset is much simple and and we hope the bloat in  
>> struct page will be compensated by the benefits in memory controllers  
>> in terms of performance and simplicity.  
>>  
>> Adding too many controllers and accounting parameters to start with  
>> will make the patch too big and complex. As Balbir mentioned, we have  
>> a plan and we shall add new control parameters in stages.  
>  
> Everyone seems to have a plan ;) I don't read the containers list...  
> does everyone still have \*different\* plans, or is any sort of consensus  
> being reached?

Consensus? I believe at this point we have a sort of consensus on the  
base container infrastructure and the need for memory controller to

control RSS, pagecache, mlock, kernel memory etc. However the implementation and approach taken is still being discussed :)

--Vaidy

---

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>

---