Subject: Re: [RFC][PATCH 3/7] Data structures changes for RSS accounting
Posted by Balbir Singh on Mon, 12 Mar 2007 17:21:30 GMT
View Forum Message <> Reply to Message

On 3/12/07, Dave Hansen <hansendc@us.ibm.com> wrote:
> On Mon, 2007-03-12 at 19:16 +0300, Kirill Korotaev wrote:
> > now VE2 maps the same page. You can't determine whether this page is mapped
> > to this container or another one w/o page->container pointer.
>
> Hi Kirill,
>
> I thought we can always get from the page to the VMA. rmap provides
> this to us via page->mapping and the 'struct address_space' or anon_vma.
> Do we agree on that?
>
> We can also get from the vma to the mm very easily, via vma->vm_mm,
> right?
>
> We can also get from a task to the container quite easily.
>
> So, the only question becomes whether there is a 1:1 relationship
> between mm_structs and containers. Does each mm_struct belong to one
> and only one container? Basically, can a threaded process have
> different threads in different containers?
>
> It seems that we could bridge the gap pretty easily by either assigning
> each mm_struct to a container directly, or putting some kind of
> task-to-mm lookup. Perhaps just a list like
> mm->tasks_using_this_mm_list.
>
> Not rocket science, right?
>
> -- Dave
>
These patches are very similar to what I posted at
http://lwn.net/Articles/223829/
In my patches, the thread group leader owns the mm_struct and all
threads belong to the same container. I did not have a per container
LRU, walking the global list for reclaim was a bit slow, but otherwise
my patches did not add anything to struct page

I used rmap information to get to the VMA and then the mm_struct.
Kirill, it is possible to determine all the containers that map the
page. Please see the page_in_container() function of
http://lkml.org/lkml/2007/2/26/7.

I was also thinking of using the page table(s) to identify all pages
belonging to a container, by obtaining all the mm_structs of tasks

belonging to a container. But this approach would not work well for
the page cache controller, when we add that to our memory controller.

Balbir

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers