Subject: Re: [RFC][PATCH 3/7] Data structures changes for RSS accounting
Posted by xemul on Mon, 12 Mar 2007 17:19:13 GMT
View Forum Message <> Reply to Message

Dave Hansen 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?

Not completely. When page is unmapped from the *very last*
user its *first* toucher may already be dead. So we'll never
find out who it was.

> 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

No. The question is "how to get a container that touched the
page first" which is the same as "how to find mm_struct which
touched the page first". Obviously there's no answer on this
question unless we hold some direct page->container reference.
This may be a hash, a direct on-page pointer, or mirrored
array of pointers.

> 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.

This could work for reclamation: we scan through all the
mm_struct-s within the container and shrink its' pages, but
we can't make LRU this way.

> Not rocket science, right?
>

> -- Dave
>
> -
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
> Please read the FAQ at http://www.tux.org/lkml/
>

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