
Subject: Re: [PATCH 2/2] Make res_counter hierarchical
Posted by [Pavel Emelianov](#) on Tue, 11 Mar 2008 08:38:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

KAMEZAWA Hiroyuki wrote:

> On Tue, 11 Mar 2008 11:15:56 +0300

> Pavel Emelianov <xemul@openvz.org> wrote:

>>> Hmm? seems strange.

>>>

>>> IMO, a parent's usage is just sum of all childs'.

>>> And, historically, memory overcommit is done against "memory usage + swap".

>>>

>>> How about this ?

>>> <mem_counter_top, swap_counter_top>

>>> <mem_counter_sub, swap_counter_sub>

>>> <mem_counter_sub, swap_counter_sub>

>>> <mem_counter_sub, swap_counter_sub>

>>>

>>> mem_counter_top.usage == sum of all mem_coutner_sub.usage

>>> swap_counter_sub.usage = sum of all swap_counter_sub.usage

>> I've misprinted in y tree, sorry.

>> The correct hierarchy as I see it is

>>

> thank you.

>

>> <mem_couter_0>

>> + -- <swap_counter_0>

>> + -- <mem_counter_1>

>> | + -- <swap_counter_1>

>> | + -- <mem_counter_11>

>> | | + -- <swap_counter_11>

>> | + -- <mem_counter_12>

>> | + -- <swap_counter_12>

>> + -- <mem_counter_2>

>> | + -- <swap_counter_2>

>> | + -- <mem_counter_21>

>> | | + -- <swap_counter_21>

>> | + -- <mem_counter_22>

>> | + -- <swap_counter_22>

>> + -- <mem_counter_N>

>> + -- <swap_counter_N>

>> + -- <mem_counter_N1>

>> | + -- <swap_counter_N1>

>> + -- <mem_counter_N2>

>> + -- <swap_counter_N2>

>>

> please let me confirm.

>

> - swap_counter_X.limit can be defined independent from mem_counter_X.limit ?
> - swap_counter_N1's limit and swap_counter_N's have some relationship ?

No. The mem_counter_N_limit is the limit for all the memory, that the Nth group consumes. This includes the RSS, page cache and swap for this group and all the child groups. Since RSS and page cache are accounted together, this limit tracks the sum of (memory + swap) values over the subtree started at the given group.

> Thanks,
> -kame
>
>

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
