
Subject: Re: [PATCH] memory.min_usage (seqlock for res_counter)

Posted by [Pavel Emelianov](#) on Wed, 05 Dec 2007 09:12:22 GMT

[View Forum Message](#) <> [Reply to Message](#)

KAMEZAWA Hiroyuki wrote:

> On Tue, 04 Dec 2007 14:10:42 +0300

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

>

>> KAMEZAWA Hiroyuki wrote:

>>> This is seqlock version res_counter.

>>> Maybe this this will reduce # of spin_lock.

>>>

>>> Pavel-san, How about this ?

>> AFAIS the readlock is used only in the check_under_limit(),

>> but I think, that even if we read usage and limit values

>> in this case non-atomically, this won't result in any

>> dramatic sequence at all. No?

>>

> Reader can detect *any* changes in res_counter member which happens

> while they access res_counter between seq_begin/seq_retry.

> Memory barrier and "sequence" of seq_lock guarantees this.

> So..there is no dramatical situation.

> (it's used for accesing xtime.)

>

> I'm sorry if I miss your point.

Sorry, let me explain it in other words.

I think, that protection in reader, that guarantees that it will see the valid result, is not very important - even if we compare usage and limit not atomically nothing serious will happen (in this particular case)

> Thanks,

> -Kame

>

>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
