
Subject: Re: [RFC][PATCH] allow "unlimited" limit value.
Posted by [Pavel Emelianov](#) on Tue, 25 Sep 2007 13:34:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Balbir Singh wrote:

> Pavel Emelianov wrote:

>> Balbir Singh wrote:

>>> KAMEZAWA Hiroyuki wrote:

>>>> On Tue, 25 Sep 2007 16:19:18 +0530

>>>> Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>>>>

>>>>> Hi, Kamezawa-San,

>>>>>

>>>> Hi,

>>>>

>>>>> Your changes make sense, but not CLUI (Command Line Usage) sense.

>>>>> 1. The problem is that when we mix strings with numbers, tools that

>>>>> parse/use get confused and complicated

>>>> yes, maybe.

>>>>

>>>>> 2. ULONGLONG_MAX is a real limit, there is no such thing as unlimited.

>>>>> If the user does ever go beyond ULONGLONG_MAX, we will limit him :-)

>>>>>

>>>> Oh. res_counter.c uses LONGLONG_MAX as default value.

>>>> need fix ? or intended ?

>>> Pavel do you remember why LONG was chosen instead of ULONG?

>> To prevent the overflow in "charge" routine.

>> See, if you add two numbers less than LONG_MAX, but with

>> unsigned long type each, you will never have an overflowed result.

>>

>

> Aah.. Thanks, my memory short circuited on me.

>

>>>> And okay there is no "unlimited" state.

>>>>

>>>>> Having said that, I do wish to have a more intuitive interface for

>>>>> users. May be a perl/python script to hide away the numbers game

>>>>> from the users. What do you think?

>>>>>

>>>> I agree with you that perl/python script can hide details. but they need knowledge

>>>> about the maximum value, which is given as default value.

>>>>

>>>> In short, what I want is some value like RLIM_INFINITY in ulimit.

>>>>

>>> I like the idea of RLIM_INFINITY and how ulimit as a tool shows

>>> a value. I guess we need something like RES_COUNTER_LIMIT_MAX

>>> and the user tool can show the limit as maximum. We could also

>>> define a special number, RES_COUNTER_LIMIT_INFINITY, such that

>>> containers will not enforce limits when the limit is set to
>>> this value.
>>>
>>>> Because it seems that res_counter.c will be used for other resource control purpose,
>>>> I thought some generic way (value) to know/specify "the maximum value" is helpful for
>>>> all resource controller interface.
>>>>
>>>> If there is an consensus that treating ULONGLONG_MAX as default, it's ok.
>>>>
>>> When I worked on the first version of res_counters, I used 0 to indicate
>>> unlimited. When Pavel posted his version, I think derived from
>>> beancounters, we did not want to have unlimited containers, so he used
>>> the maximum value
>> Yup! Setting LONGMAX pages for container means that this container
>> is unlimited, since there're no such many pages on any arch :)
>>
>
> Pavel, we no longer account in pages, we do it in bytes. Second,
> this assumption cannot hold for long, memory sizes are growing,
> I think we need a special value.

I see. And I also see that we've switched into unsigned long long.

Well, no container may have the ULLMAX (or what is it?) bytes touched/allocated :) So I don't see any need in a special value.

>
>>>> Thanks,
>>>> -Kame
>>>>
>>> Thanks for looking into this,
>>>
>
>

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