Subject: Re: [PATCH] Simplify memory controller and resource counter I/O
Posted by Balbir Singh on Fri, 05 Oct 2007 03:45:46 GMT
View Forum Message <> Reply to Message

Paul Menage wrote:
> On 10/4/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
>> Forbidding writing to the root resource counter is a policy decision
>> I am unable to make up my mind about. It sounds right, but unless
>> we have a notion of unlimited resources, I am a bit concerned about
>> taking away this flexibility.
>
> One big reason for doing this is to make virtualization easier - if
> you expect not to be able to write to your root cgroup's limits files,
> then it's easier to make them non-writeable for a virtual server.
>

Can't we handle that through file system permissions? virtual servers
will not run as root

>>>> One arguable drawback to this patch is that the use of memparse() is
>>>> lost in the cleanup. Having said that, given the existing of shell
>>>> arithmetic, it's not clear to me that typing
>>>>
>> memparse(), makes it so much easier, we need to use it.
>>
>>>> echo $[2<<30] > memory.limit
>>>>
>> Very geeky! I don't like it personally
>
> Why do you dislike it? Do you really believe that anyone using this
> interface by hand isn't going to know that MB is 2^20 and GB is 2^30?
>

But system administrators deal with memory in MB and GB. When you go
to buy memory, you don't specify, I need 1 << 30 or 2^30 bytes of
memory :-). Most administrators track their memory using these
quantifiers.

>> Do read_uint() and write_uint(), just read and write unsigned
>> integers?
>
> Correct.
>

Oops.. that would be problem, what if I wanted to set my limit to
unsigned long long max?

> Paul

--
Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

_____

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