
Subject: Re: [RFC][PATCH 1/7] Resource counters
Posted by [ebiederm](#) on Sun, 11 Mar 2007 19:00:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Poetzl <herbert@13thfloor.at> writes:

>
> Linux-VServer does the accounting with atomic counters,
> so that works quite fine, just do the checks at the
> beginning of whatever resource allocation and the
> accounting once the resource is acquired ...

Atomic operations versus locks is only a granularity thing.
You still need the cache line which is the cost on SMP.

Are you using `atomic_add_return` or `atomic_add_unless` or
are you performing you actions in two separate steps which
is racy? What I have seen indicates you are using a racy two separate
operation form.

>> If we'll remove `failcnt` this would look like
>> `while (atomic_cmpxchg(...))`
>> which is also not that good.
>>
>> Moreover - in RSS accounting patches I perform page list
>> manipulations under this lock, so this also saves one atomic op.
>
> it still hasn't been shown that this kind of RSS limit
> doesn't add big time overhead to normal operations
> (inside and outside of such a resource container)
>
> note that the 'usual' memory accounting is much more
> lightweight and serves similar purposes ...

Perhaps....

Eric

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