
Subject: Re: Q: Do systems using containers user more process ids?

Posted by [ebiederm](#) on Mon, 14 Aug 2006 22:07:02 GMT

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

Dave Hansen <haveblue@us.ibm.com> writes:

> On Mon, 2006-08-14 at 15:01 -0600, Eric W. Biederman wrote:

>> The practical question is if systems using containers are using noticeably
>> more pids than anyone else. So far the responses I have gotten indicate
>> that users aren't. So at least until we descend into multi-core madness
>> it sounds like the current structures are fine, but it might be worth moving
>> the cap on the number of pid hash table entries at some point in the future.
>

> Since it is already resized at boot-time, I can't imagine this be a real
> problem to fix. I assume you're just trying to see if anybody has run
> into it as of yet.

More or less. There is some other work that needs to be done and I'm
trying to see if reworking the data structure make sense. Not having
a hash table would be nice in the container case.

> Perhaps a one-time-per-boot warning in find_pid() if the chains get too
> long would be nice to have. It wouldn't give us detailed performance
> measurements, but it would be a nice canary in the mine in case
> something goes horribly wrong.
>
> What about something like this?

If we put the canary on pushing pid_max up I'm all for it.
Our current hash is even enough and our set of pids small enough
that just knowing pid_max we can easily calculate the worst case hash
chain length, by brute force.

Eric
