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

Posted by [dev](#) on Tue, 15 Aug 2006 08:13:03 GMT

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

>>We have not seen any degradation here in real cases,
>>but probably you are right and pid hash can be allocated taking into account
>>physical memory as it is done for TCP/ip/other hashes?
>
>
> It is but it is currently capped at 4K entries.
> With 4K entries and 32K pids our worst case is usage is a hash chain
> 9 entries long. At 4M pids our hash chains are 1000 entries long, which
> sucks.
4M pids are almost unreal in production systems.
(only if you spawn these tasks to make them sleep forever :)))).
we usually have no more than 20,000 tasks (which is 200VEs with 100 tasks in each)

>>But not sure, it is worth bothering right now... Maybe it worth first to make
>>some
>>simple test, say:
>>
>>1. run 50,000 tasks.
>>2. run some benchmark
>>
>>and compare benchmark results with different hash sizes?
>>What do you think?
>
>
> If it is easy sure. The real point of where things degrade is
> past 50K processes though.
>
> 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.
containers are using noticeably more pids, I think it is not a doubt...
the question is whether it is worth doing something here _now_...

Kirill
