
Subject: Namespaces exhausted CLONE_XXX bits problem
Posted by [Pavel Emelianov](#) on Mon, 14 Jan 2008 13:45:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi, guys!

I started looking at PTYs/TTYs/Console to make the appropriate namespace and suddenly remembered that we have already exhausted all the CLONE_ bits in 32-bit mask.

So, I recalled the discussions we had and saw the following proposals of how to track this problem (with their disadvantages):

1. make the clone2 system call with 64-bit mask
 - this is a new system call
2. re-use CLONE_STOPPED
 - this will give us only one bit
3. merge existing bits into one
 - we lose the ability to create them separately
4. implement a sys_unshare_ns system call with 64bit/arbitrary mask
 - this is anew system call
 - this will bring some dissymmetry between namespaces
5. use sys_indirect
 - this one is not in even -mm tree yet and it's questionable whether it will be at all

I have one more suggestion:

6. re-use bits, that don't make sense in sys_unshare (e.g. CLONE_STOPPED, CLONE_PARENT_SETTID, CLONE_VFORK etc)
This will give us ~16 new bits, but this will look not very nice.

What do you think about all of this?

Thanks,
Pavel

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