

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](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---