

---

Subject: Re: userns: targeted capabilities v5  
Posted by [ebiederm](#) on Fri, 18 Feb 2011 03:53:37 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Andrew Morton <[akpm@linux-foundation.org](mailto:akpm@linux-foundation.org)> writes:

> On Thu, 17 Feb 2011 15:02:24 +0000  
> "Serge E. Hallyn" <[serge@hallyn.com](mailto:serge@hallyn.com)> wrote:  
>  
>> Here is a repost of my previous user namespace patch, ported onto  
>> last night's git head.  
>>  
>> It fixes several things I was doing wrong in the last (v4)  
>> posting, in particular:  
>>  
>> 1. don't set uts\_ns->user\_ns to current's when !CLONE\_NEWUTS  
>> 2. add a ipc\_ns->user\_ns which owns ipc\_ns, and use that to  
>> decide CAP\_IPC\_OWNER  
>> 3. fix logic flaw caused by bad parantheses  
>> 4. allow do\_prlimit to current  
>> 5. don't always give root full privs to init\_user\_ns  
>>  
>> The expected course of development for user namespaces is laid out  
>> at <https://wiki.ubuntu.com/UserNamespace>.  
>  
> Seems like a nice feature to be developing.  
>  
> I worry about the maturity of it all at this stage. How far along is  
> it \*really\*?  
>  
> Is anyone else working with you on developing and reviewing this work?

I did a lot of the initial design and I have been reviewing as I have time.

Andrew at a practical level we have to merge this incrementally. Anything much bigger than Serge's current patchset will be too big to review. The first really bit step is making capabilities local to the user namespace, and that is what this patchset does along with using that localness in some good places.

Eric

---

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

---