

---

Subject: Re: [PATCH 0/4] user namespaces: introduction

Posted by [Daniel Hokka Zakrisson](#) on Mon, 28 Jan 2008 19:24:24 GMT

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

---

Serge E. Hallyn wrote:

> Here is a small patchset I've been sitting on for awhile  
> to make signaling mostly subject to user namespaces. In  
> particular,  
>  
> 1. store user\_namespace in user struct  
> 2. introduce CAP\_NS\_OVERRIDE  
> 3. require CAP\_NS\_OVERRIDE to signal another user namespace  
>  
> The first step should have been done all along. Else wouldn't  
> a hash collision on (ns1, uid) and (ns2, uid), however unlikely,  
> give us wrong results at uid\_hash\_find()?

Unless I've completely misunderstood the code, each namespace has a separate hash. Please correct me if I'm wrong.

> The main remaining signaling+usersns issue is of course the  
> siginfo. Tacking a usersns onto siginfo is a pain due to  
> lifetime mgmt issues. I haven't decided whether to just  
> catch all the callers and fake uid=0 if user namespaces  
> aren't the same, introduce some unique non-refcounted id to  
> represent (user,user\_ns), or find some other way to deal with  
> it.  
>  
> thanks,  
> -serge

--

Daniel Hokka Zakrisson

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---