
Subject: Re: [RFC][PATCH 0/2] user namespace [try #2]
Posted by [Cedric Le Goater](#) on Mon, 11 Sep 2006 08:09:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Herbert Poetzl <herbert@13thfloor.at> writes:

>

>> On Thu, Sep 07, 2006 at 12:18:14PM -0600, Eric W. Biederman wrote:

>>> Kirill Korotaev <dev@sw.ru> writes:

>>>

>>>> yes, these patches are usable for OpenVZ AS IS, so I'm not sure

>>>> why we can't do step by step and commit. However I posted some comments on

>>>> patches...

>>>>

>>>> Eric do you have some STRONG objections (maybe I just missed it somewhere)?

>>> - We do not handle interactions between processes in different uid

>>> namespaces and still have the normal uid equality checks.

>>> - I am willing to be convinced that this is a nuclear missile the user

>>> is allowed to shoot themselves in the foot with if someone can show me

>>> how to use the current version safely.

>>>

>>> A lot of this scares me silly as when ever you touch the primary

>>> identifier in the security checks you must be very very very careful.

>>> My gut feeling is that I'm nowhere near paranoid enough and the rest

>>> of you aren't even paranoid.

>>>

>>> What I want to see is that every uid identity check becomes either

>>> a struct user comparison or a uid, uid_ns tuple comparison.

>> second that!

>

> In addition I don't have problems with incremental progress

> if we implement in such a way that we don't enable the ability

> to create a new uid namespace to user space before we are certain

> it is safe.

>

> All of the code could be present and we just have a one line check

> that denied requests to create a new namespace.

OK. I'll see how this is possible. I guess the simplest way for the moment is to remove the unshare() of the user_namespace.

So, shall we follow the 'grep' method for uids like we are doing for pids and thread ? This is going to be painful but I guess there is no simple solution ...

C.

Containers mailing list

