Subject: Re: [patch 0/8] unprivileged mount syscall Posted by Miklos Szeredi on Fri, 13 Apr 2007 11:58:59 GMT View Forum Message <> Reply to Message

```
> On Wed, 2007-04-11 at 12:44 +0200, Miklos Szeredi wrote:
>>> 1. clone the master namespace.
>>> 2. in the new namespace
>>> move the tree under /share/$me to /
>>>
          for each ($user, $what, $how) {
             move /share/$user/$what to /$what
>>>
        if (\text{show} == \text{slave}) {
>>>
               make the mount tree under /$what as slave
>>>
>>>
>>>
          }
>>>
>> 3. in the new namespace make the tree under
         /share as private and unmount /share
> >
>> Thanks. I get the basic idea now: the namespace itself need not be
>> shared between the sessions, it is enough if "share" propagation is
> > set up between the different namespaces of a user.
> >
>> I don't yet see either in your or Viro's description how the trees
> > under /share/$USER are initialized. I guess they are recursively
> > bound from /, and are made slaves.
> yes. I suppose, when a userid is created one of the steps would be
> mount --rbind / /share/$USER
> mount --make-rslave /share/$USER
> mount --make-rshared /share/$USER
```

Thinking a bit more about this, I'm quite sure most users wouldn't even want private namespaces. It would be enough to

chroot /share/\$USER

and be done with it.

Private namespaces are only good for keeping a bunch of mounts referenced by a group of processes. But my guess is, that the natural behavior for users is to see a persistent set of mounts.

If for example they mount something on a remote machine, then log out from the ssh session and later log back in, they would want to see their previous mount still there.

Miklos

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