Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag
Posted by ebiederm on Mon, 16 Apr 2007 19:16:48 GMT
View Forum Message <> Reply to Message

Miklos Szeredi <miklos@szeredi.hu> writes:

>> > That depends. Current patches check the "unprivileged submounts
>> > allowed under this mount" flag only on the requested mount and not on
>> > the propagated mounts. Do you see a problem with this?
>>
>> I think privileges of this sort should propagate. If I read what you
>> just said correctly if I have a private mount namespace I won't be able
>> to mount anything unless when it was setup the unprivileged submount
>> command was explicitly set.
>
> By design yes. Why is that a problem?

It certainly doesn't match my intuition.

Why are directory permissions not sufficient to allow/deny non-priveleged mounts?
I don't understand that contention yet.

I should probably go back and look and see how plan9 handles mount/unmount
permissions. Plan9 gets away with a lot more because it doesn't have
a suid bit and mount namespaces were always present, so they don't have
backwards compatibility problems.

My best guess at the moment is that plan9 treated mount/unmount as
completely unprivileged and used the mount namespaces to limit the
scope of what would be affected by a mount/unmount operation. I think
that may be reasonable in linux as well but it will require the
presence of a mount namespace to limit the affects of what a user can
do.

So short of a more thorough audit I believe the final semantics should
be:
- mount/unmount for non-priveleged processes should only be limited
  by the mount namespace and directory permissions.
- CLONE_NEWNS should not be a privileged operation.

What prevents us from allowing these things?

- Unprivileged CLONE_NEWNS and unprivileged mounts needs resource
  accounting so we don't have a denial of service attack.

- Unprivileged mounts must be limited to directories that we have
  permission to modify in a way that we could get the same effect

as the mount or unmount operation in terms of what files are visible
otherwise we can mess up SUID executables.

- Anything else?

There are user space issues such as a reasonable pam module and how
to do backups. However those are user space issues.

What am I missing that requires us to add MNT_USER and MNT_USERMNT?

Eric

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