Subject: Re: namespaces compatibility list
Posted by ebiederm on Tue, 06 Nov 2007 17:46:23 GMT
View Forum Message <> Reply to Message

Pavel Emelyanov <xemul@openvz.org> writes:

> Eric W. Biederman wrote:
>> Cedric Le Goater <clg@fr.ibm.com> writes:
>>> right. I think we can address Ulrich concerns first because we have
>>> a solution for it (which looks like unsharing all namespaces at once,
>>> here comes back the container object story :)
>>
>> It doesn't work because we can't create a fresh mount namespace.
>>
>> We need to create all new mounts (and deny access to the old ones)
>> if we want to prevent all possibility of user space goof ups.
>>
>> While that is easy enough to build an application to do we can't
>> easily enforce that in the kernel. Currently this is all
>> CAP_SYS_ADMIN so only root can do this anyway. So we can easily
>> say don't do that then.
>>
>> Clone flag consistency checking should only be used to enforce
>> cases where the kernel side cannot support correctly. Currently
>> the kernel has no problems with the current mix and match possibilities
>> short of implementation deficiencies. So I do not see us
>> addressing Ulrich's concerns with clone flags.
>
> ACK :) Since this all is CAP_SYS_ADMIN-ed we can do with just a warning.

So to restate.

clone flags consistency checks are for things the kernel can't do or
  for things that the kernel can't do securely.

If all we do is confuse user space if used improperly it's simply
a don't do that then.

CAP_SYS_ADMIN keeps us untrusted applications from confusing suid
executables, which is the only case where confusion counts as a
security hole.

Eric

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