Subject: Re: [patch -mm 08/17] nsproxy: add hashtable
Posted by Cedric Le Goater on Tue, 12 Dec 2006 18:29:08 GMT
View Forum Message <> Reply to Message

>>> If someones permissions to various objects does not depend on the namespace
>>> they are in quite possibly this is a non-issue. If we actually depend on
>>> the isolation to keep things secure enter is a setup for a first rate escape.
>> I don't believe the isolation can be effective between two namespaces
>> where one is an ancestor of another. It can be so long as one isn't
>> the ancestor of another, but then we're not allowing either to enter
>> the other namespace. So it's not a problem.
>
> Reasonable.
>
>> The bind_ns() proposed by Cedric is stricter, only allowing nsid 0 to
>> switch namespaces. So it may be overly restrictive, and does introduce
>> a new global namespace, but it is safe.
>
> I will look a little more. There are a lot patches out there that need
> review. What disturbs a little is that with ptrace we have an existing
> mechanism that can do everything we want enter or bind_ns to be able to do.

Eric, you have this habit of flooding us with email whenever a patchset is
sent on this topic. It is a bad habit. Please take some time to look at it
before. There is work behind it and it tries to address some issues.

This patchset has been sent on container@ as a proposal for -mm. I'll try
to make a summary of how we can improve next one to move forward.

I still need to read all your emails :)

thanks,

C.

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers