Subject: Re: [patch 1/2] [RFC] Simple tamper-proof device filesystem.
Posted by Oren Laadan on Tue, 18 Dec 2007 03:03:42 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> Quoting Oren Laadan (orenl@cs.columbia.edu):
>> I hate to bring this again, but what if the admin in the container
>> mounts an external file system (eg. nfs, usb, loop mount from a file,
>> or via fuse), and that file system already has a device that we would
>> like to ban inside that container ?
>
> Miklos' user mount patches enforced that if !capable(CAP_MKNOD),
> then mnt->mnt_flags |= MNT_NODEV. So that's no problem.

Yes, that works to disallow all device files from a mounted file system.

But it's a black and white thing: either they are all banned or allowed;
you can't have some devices allowed and others not, depending on type
A scenario where this may be useful is, for instance, if we some apps in
the container to execute withing a pre-made chroot (sub)tree within that
container.

>
> But that's been pulled out of -mm! ? Crap.
>
>> Since anyway we will have to keep a white- (or black-) list of devices
>> that are permitted in a container, and that list may change even change
>> per container -- why not enforce the access control at the VFS layer ?
>> It's safer in the long run.
>
> By that you mean more along the lines of Pavel's patch than my whitelist
> LSM, or you actually mean Tetsuo's filesystem (i assume you don't mean that
> by 'vfs layer' :), or something different entirely?

:)

By 'vfs' I mean at open() time, and not at mount(), or mknod() time.
Either yours or Pavel's; I tend to prefer not to use LSM as it may
collide with future security modules.

Oren.

>
> thanks,
> -serge

_____
Containers mailing list
Containers@lists.linux-foundation.org