Subject: Re: [PATCH] cgroups: implement device whitelist lsm (v3)
Posted by serue on Fri, 14 Mar 2008 14:32:16 GMT
View Forum Message <> Reply to Message

Quoting Stephen Smalley (sds@epoch.ncsc.mil):
>
> On Fri, 2008-03-14 at 21:17 +1100, James Morris wrote:
> > On Thu, 13 Mar 2008, Serge E. Hallyn wrote:
> >
> > > Implement a cgroup using the LSM interface to enforce open and mknod
> > > on device files.
> >
> > Actually, I'm not sure that the LSM approach in general is best here.
> >
> > The LSM model is that standard DAC logic lives in the core kernel, and
> > that extended security logic (e.g. MAC) is called after DAC via hooks.
> > cgroups has introduced new security logic of its own, which is arguably
> > "standard DAC" when cgroups is enabled.
> >
> > I can understand Greg not wanting this security logic in the core kernel,
> > but it is specific to cgroups (which itself is security model agnostic)
> > and does not stand alone as a distinct security framework.

I completely disagree.   We have two separate frameworks in the kernel,
one to enforce generic additional security stuff, and one to track
tasks. When I need a feature which tracks tasks to do some security
tasks, it seems obvious that I would use both, just like to enforce a
certain type of MAC I end up using both netfilter and LSM through
selinux.

> > The fact that all existing LSMs need to invoke exactly the same code is an
> > indicator that it doesn't belong in LSM.

No, that's like saying capabilities don't belong in LSM because all LSMS
need to invoke it the same way. What it is an indicator of is that
there are (not-quite-)orthogonal pieces of security which users might
want to use together.

As I told stephen I hope to provide the enhanced selinux support for
devices, and at that point perhaps you won't want to support
SELINUX+CGROUPS_DEV anymore.

Now that's just my opinion and it doesn't count for much. I'll do
whatever everyone can agree on, but will wait for Paul's opinion about
adding cgroup hooks next to the two security hooks.

> > Moving this logic into LSM means that instead of the cgroups security
> > logic being called from one place in the main kernel (where cgroups

> > lives), it must be called identically from each LSM (none of which are
> > even aware of cgroups), which I think is pretty obviously the wrong
> > solution.
> >
> > This is baggage which comes with cgroups -- please don't push it into LSM
> > to try and hide that.
>
> I agree with the above, and would further note that I would expect the
> SELinux solution to the problem would be done not by stacking with or
> calling this device whitelist lsm but instead by introducing the ability
> to bind security labels to devices within the kernel (independent of the
> particular device node(s) in the filesystem used to access that device)
> and applying permission checks on those device labels when processes
> attempt to create or access those devices (again independent of the
> particular device node used to access them). That keeps the policy
> integrated and analyzable and avoids an external dependency.

Agreed.

-serge

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