Subject: Re: [PATCH] cgroups: implement device whitelist lsm (v3)
Posted by James Morris on Fri, 14 Mar 2008 10:17:12 GMT
View Forum Message <> Reply to Message

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.

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

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.

- James
--
James Morris
<jmorris@namei.org>

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