Subject: Re: LSM and Containers
Posted by Peter Dolding on Wed, 24 Oct 2007 23:07:52 GMT
View Forum Message <> Reply to Message

The other thing you have not though of and is critical. If LSM is the same LSM across all containers. What happens if that is breached and tripped to disable. You only want to loss one container to a breach not the whole box and dice in one hit. Its also the reason why my design does not have a direct link between controllers. No cascade threw system to take box and dice.

The more I look at it more holes I find why the current LSM model just cannot keep on existing with Containers. Its not the best option. Hacking it to work with containers is only creating risks of more problems. The LSM model as also breed that problem of not sharing security tech advantages to everyone. Ie if they don't use our LSM they don't need/deserve our defense.

Different LSM per container from a security point of view appears critical. Sorry to say redesign from the ground up time everyone. Its a round peg into a square hole yes you can bash it in but it will never fit right.

Peter Dolding

ps sorry for going on so long I just see this as a major problem. If you have a solution to it tell me. Since a cut line has be put somewhere with containers.

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

Subject: Re: LSM and Containers
Posted by Crispin Cowan on Wed, 24 Oct 2007 23:21:44 GMT
View Forum Message <> Reply to Message

Peter Dolding wrote:
> The other thing you have not though of and is critical. If LSM is the
> same LSM across all containers. What happens if that is breached and
> tripped to disable. You only want to loss one container to a breach
> not the whole box and dice in one hit. Its also the reason why my
> design does not have a direct link between controllers. No cascade
> threw system to take box and dice.
>
Sorry, but I totally disagree.

If you obtain enough privilege to disable the LSM in one container, you
also obtain enough privilege to disable *other* LSMs that might be
operating in different containers. This is a limitation of the
Containers feature, not of LSM.

The purpose of LSM would be to manage privilege such that you cannot do
damage, and in particular, any LSM that fails to prevent an attacker
from disabling the LSM itself has failed, either in design, or in having
an inadequate policy in place.

> The more I look at it more holes I find why the current LSM model just
> cannot keep on existing with Containers. Its not the best option.
> Hacking it to work with containers is only creating risks of more
> problems. The LSM model as also breed that problem of not sharing
> security tech advantages to everyone. Ie if they don't use our LSM
> they don't need/deserve our defense.
>
Again, I completely disagree.

Well, I agree that the hacking you proposed to permit different LSMs in
different containers is a bad idea, so lets not do that :)

I see no need to support different LSMs in different containers. The
complexity of such a feature would be very high. The utility strikes me
as being very low; people who want that degree of separation of
containers should be using Xen or KVM, not Containers.

> Different LSM per container from a security point of view appears
> critical. Sorry to say redesign from the ground up time everyone.
> Its a round peg into a square hole yes you can bash it in but it will
> never fit right.
>
I have no idea how you can support such assertions. Absolutely not. It
is quite clear that the way to address security for containers is to
enhance individual LSM modules to be container-aware so that you can
have separate policies in the separate containers. That is in keeping
with the spirit of sharing the kernel, and providing separate instances
to the users.

> ps sorry for going on so long I just see this as a major problem. If
> you have a solution to it tell me. Since a cut line has be put
> somewhere with containers.
>
Where as I see it as a very minor problem, and very easy to fix without
any re-design of LSM, or of Containers. It only requires container-aware
LSM modules.

Crispin

--
Crispin Cowan, Ph.D. http://crispincowan.com/~crispin/
Itanium. Vista. GPLv3. Complexity at work

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

---

Subject: Re: LSM and Containers
Posted by Peter Dolding on Thu, 25 Oct 2007 00:20:37 GMT
View Forum Message <> Reply to Message

On 10/25/07, Crispin Cowan <crispin@crispincowan.com> wrote:
> Peter Dolding wrote:
> > The other thing you have not though of and is critical. If LSM is the
> > same LSM across all containers. What happens if that is breached and
> > tripped to disable. You only want to loss one container to a breach
> > not the whole box and dice in one hit. Its also the reason why my
> > design does not have a direct link between controllers. No cascade
> > threw system to take box and dice.
> >
> Sorry, but I totally disagree.
>
> If you obtain enough privilege to disable the LSM in one container, you
> also obtain enough privilege to disable *other* LSMs that might be
> operating in different containers. This is a limitation of the
> Containers feature, not of LSM.
>
That is not a Container feature. If you have enough privilege does
not mean you can. Root user in a Container does not mean you can play
with other containers applications. There is a security split at the
container edge when doing Virtual Servers what by using one LSM you
are disregarding.

Simple point if one LSM is disabled in a container it can only get the
max rights of that Container. So cannot see the other LSM's on the
system bellow it. Reason also why in my model its the same layout if
there is 1 or 1000 stacked so attack cannot tell how deep they are in
and if there is anything to be gained by digging. You have to break
the Security Container as well to get higher. Of course breaking the
base LSM you would have problems the one with full powers of the
system and the right to kill and control the containers bellow it.
Most likely it would be wise to run that just for limited operations.
Really limited operations.

I think you need to go any play with Solaris some time what containers can do is quite impressive. Right upto changing the syscalls and device names inside them. Now its only going to get tricker if someone brings in Freebsd and Solaris emulation containers into linux. Yes this is still just using the Linux kernel. Because LSM's don't exist on them you will want to plug in a module to suit the platform contained inside the container. In particular something different to process the security configs even if the same enforcement code is used.

You are dealing with something far more powerful current model is not going to fit. I look long term and I just cannot find what you are doing is going to fit in anyway shape or form. Just different Linux distros is the simple form of containers not is fully grown form. This is the problem if you cannot do that future of containers will have problems because LSM will be limiting its forms.

To rebuild a security framework takes time. Its time to pull head out sand that containers are not simple thing that what has worked in past with small alterations will work with it in future.

Containers are completely new system with completely new problems. LSM's should not interfere with its development.

What you are doing is having each LSM create its own outer shield. With its own weaknesses. Common Security Container makes only one shield at edge of containers to maintain and look after. Does not different weakness to track down and fix on a LSM by LSM base.

Peter Dolding

_____

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