

---

Subject: Re: [RFC][patch 0/3] Network container subsystem - bind filtering  
Posted by [serue](#) on Wed, 05 Sep 2007 15:37:41 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Quoting dlezcano@fr.ibm.com (dlezcano@fr.ibm.com):

> Paul Menage mentionned, a few weeks ago, he wanted a bind filtering  
> for containers. Here it is :)  
>  
> The following patches are a proposition to bring IP isolation to a container.  
>  
> After looking more closely at the code I found that security hooks are  
> at the right place to catch socket calls. The IP isolation relies on the  
> security hooks and that has the advantage of not having the kernel code modified,  
> (expect container.h and makefile/kconfig), the patchset provide just a new  
> file container\_network.c  
>  
> Roughly, a container has a subsystem for the network (only ipv4).

Just curious - why ipv4 only? When i wrote the bsdjail lsm, using the same approach, doing ipv6 address was pretty simple. Did something change? Or is ipv4 just a temporary restriction while you prototype?

> This subsystem contains the list of the addresses allowed to be used by the  
> container. If a container tries to bind to an address not contained into  
> this list, the bind will fail with EPERM. Of course the bind is allowed to  
> INADDR\_ANY.  
>  
> If this approach is ok for everyone, I can extend the bind filtering to  
> consolidate the IP isolation.

You'll want to cc: the linux-security-module@vger.kernel.org list on this patchset.

thanks,  
-serge

---

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

---