
Subject: Re: [RFC][patch 0/3] Network container subsystem - bind filtering
Posted by [Benjamin Thery](#) on Wed, 05 Sep 2007 14:05:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

The patchset looks fine to me.
This looks like a nice and simple way to obtain basic IP isolation.

Benjamin

dlezcano@fr.ibm.com wrote:

> 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,
> (except 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).
> 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.
>
> Regards.
>

--
Benjamin Thery - BULL/DT/Open Software R&D

<http://www.bull.com>

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