
Subject: Re: [PATCH 12/16] net: Support multiple network namespaces with netlink
Posted by [Pavel Emelianov](#) on Mon, 10 Sep 2007 13:46:48 GMT

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

Eric W. Biederman wrote:

> Each netlink socket will live in exactly one network namespace,
> this includes the controlling kernel sockets.
>
> This patch updates all of the existing netlink protocols
> to only support the initial network namespace. Request
> by clients in other namespaces will get -ECONREFUSED.
> As they would if the kernel did not have the support for
> that netlink protocol compiled in.
>
> As each netlink protocol is updated to be multiple network
> namespace safe it can register multiple kernel sockets
> to acquire a presence in the rest of the network namespaces.
>
> The implementation in af_netlink is a simple filter implementation
> at hash table insertion and hash table look up time.
>
> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

> ---

> drivers/connector/connector.c | 2 +-
> drivers/scsi/scsi_netlink.c | 2 +-
> drivers/scsi/scsi_transport_iscsi.c | 2 +-
> fs/ecryptfs/netlink.c | 2 +-
> include/linux/netlink.h | 6 ++-
> kernel/audit.c | 4 +-
> lib/kobject_uevent.c | 5 +-
> net/bridge/netfilter/ebt_ulog.c | 5 +-
> net/core/rtnetlink.c | 4 +-
> net/decnnet/netfilter/dn_rtmsg.c | 3 +-
> net/ipv4/fib_frontend.c | 4 +-
> net/ipv4/inet_diag.c | 4 +-
> net/ipv4/netfilter/ip_queue.c | 6 +-
> net/ipv4/netfilter/ipt_ULOG.c | 3 +-
> net/ipv6/netfilter/ip6_queue.c | 6 +-
> net/netfilter/nfnetlink.c | 2 +-
> net/netfilter/nfnetlink_log.c | 3 +-
> net/netfilter/nfnetlink_queue.c | 3 +-
> net/netlink/af_netlink.c | 106 ++++++-----
> net/netlink/genetlink.c | 4 +-
> net/xfrm/xfrm_user.c | 2 +-
> security/selinux/netlink.c | 5 +-
> 22 files changed, 122 insertions(+), 61 deletions(-)

Rrrrrr. This is the 5th or even the 6th patch that changes tens of files

but (!) most of these changes are just propagating some core thing into protocols, drivers, etc. E.g. you add an argument to some function and then make all the rest use it, but the chunk adding the argument itself is buried in these changes.

Why not make a reviewers' lives easier and make (with hands) the core hunks go first and the "propagation" ones at the end? For RFC purpose I would even break the git-bisect safeness and splitted these patches into 2 parts: those with the core and those with the propagation.

Thanks,
Pavel

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