

Re: [RFC][PATCH 0/9] Network receive deadlock prevention for NBD

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-08/msg04718.html>

- *From:* Peter Zijlstra <a.p.zijlstra@xxxxxxxxx>
 - *Date:* Mon, 14 Aug 2006 08:45:43 +0200
-

On Mon, 2006-08-14 at 09:13 +0400, Evgeniy Polyakov wrote:

On Sun, Aug 13, 2006 at 01:16:15PM -0700, Daniel Phillips (phillips@xxxxxxxxx) wrote:

Indeed. The rest of the corner cases like netfilter, layered protocol and so on need to be handled, however they do not need to be handled right now in order to make remote storage on a lan work properly. The sane thing for the immediate future is to flag each socket as safe for remote block IO or not, then gradually widen the scope of what is safe. We need to set up an opt in strategy for network block IO that views such network subsystems as ipfilter as not safe by default, until somebody puts in the work to make them safe.

Just for clarification – it will be completely impossible to login using openssh or some other priveledge separation protocol to the machine due to the nature of unix sockets. So you will be unable to manage your storage system just because it is in OOM – it is not what is expected from reliable system.

But really, if you expect to run reliable block IO to Zanzibar over an ssh tunnel through a firewall, then you might also consider taking up bungie jumping with the cord tied to your neck.

Just pure openssh for control connection (admin should be able to login).

These periods of degenerated functionality should be short and infrequent albeit critical for machine recovery. Would you rather have a slower ssh login (the machine will recover) or drive/fly to Zanzibar to physically reboot the machine?

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

Re: [RFC][PATCH 0/9] Network receive deadlock prevention for NBD

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>