
Subject: Re: [PATCH] ipv4/ipvs: Convert to kthread API
Posted by [Christoph Hellwig](#) on Sun, 22 Apr 2007 19:50:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Apr 19, 2007 at 03:59:44PM -0700, Andrew Morton wrote:
> There still seems to be quite a lot of complexity in this driver's
> thread handling which could be removed if we did a full conversion
> to the kthread API.
>
> It all looks.... surprisingly complex in there.

It is. There quite a few interesting oddities in this code:

- creation of a forker thread. This is superfluous when using the kthread infrastructure as a thread created by `kthread_create` always comes from our dedicated forker thread.
- the infinite retry on failure looks very bogus, the system doesn't recover very well if you try to fork forever in a loop :)
- a lot of very overlapping state variables. My reading of the code suggests that both a 'master' and 'backup' thread can run at the same time. I think the code would benefit a lot from totally separating these codepaths.
- `start_sync_thread` and `stop_sync_thread` are called with unchecked user supplied arguments and bug if they don't match the expected values. While all this is under `capable(CAP_NET_ADMIN)` it still sounds like something to fix.
- and the usual removal of semaphores and completions for startup/shutdown would benefit the code a lot, as for most thread users.

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