
Subject: Re: [PATCH 2/3] [NETNS49] Add struct net to all calls for rt_cache_flush
Posted by [ebiederm](#) on Fri, 19 Oct 2007 19:23:42 GMT

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

"Denis V. Lunev" <den@sw.ru> writes:

> I think that this is quite boring. It is impossible to get the namespace almost
> anywhere even when it is safe to use this :(

Well for sysctl is my intention to eventually have:

/proc/sys -> /proc/self/sys

/proc/<pid>/sys

So it doesn't depend on who you are so much as who you are doing the work for. Especially in the networking context.

The goal for filesystem interfaces to namespaces is that long term we should not care who is accessing the file but rather which file is being accessed. This allows for monitoring and control applications outside the container. Applications like the planetlab monitoring framework where the monitoring applications runs in a low privilege container and can still monitor the entire box.

I don't quite have the user interface side of the network namespace proc and sysctl interfaces sorted out but I do have the internal interfaces working in a way that supports this.

> extra2 is a bad place, as it has a conventional meaning for other kernel. For
> example it is used for quite a lot of sysctl-s anywhere.

extra1 and extra2 are defined to be whatever the user of that sysctl table entry want them to be. Frequently they are used for min/max behavior but if your sysctl has another use for them that is fine.

Eric
