
Subject: Re: [PATCH] [NETNS49] support for per/namespace routing cache cleanup

Posted by [Benjamin Thery](#) on Thu, 18 Oct 2007 07:18:48 GMT

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

Denis V. Lunev wrote:

> Daniel Lezcano wrote:

>> Oh, by the way, I forgot something important you spotted with the list

>> protected by the mutex.

>>

>> When looking at ipv6/fib_hash.c with Benjamin, we need to browse the

>> network namespaces list for the garbage collecting, but we are in an

>> interrupt handler, so I can not use rtnl_lock.

> where exactly....

Actually, it is in net/ipv6/ip6_fib.c, in fib6_clean_all().

fib6_clean_all() is called by fib6_run_gc() handler of the ip6_fib_timer.

If we don't want to have one such timer per net namespace, in fib6_clean_all() we have to go through all net to clean their own fib_table_hash (using for_each_net() protected by rtnl_lock).

> all interesting places are called under rtnl already...

>

>> Why is not possible to protect the list with a simple spinlock ? so we

>> can call spin_lock_bh when we are in interrupt handler.

>

> see my old patch with a locking rules :)

>

> By the way, I have forgotten to mention, that the original patch works

> from 2004 for OpenVz :)

>

> Regards,

> Den

>

--

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

<http://www.bull.com>
