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

Posted by [Daniel Lezcano](#) on Fri, 19 Oct 2007 07:39:24 GMT

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

Denis V. Lunev wrote:

> Benjamin Thery wrote:

>> 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).

>

> after careful thinking, one timer per/namespace looks better for me :)

So you prefer to have multiple timers flushing a portion of a shared routing cache ?

It is not impossible I already tried that for ipv6. it will just touch much more files.

If we want, at all costs, to have the routing cache flushed per namespace, I am much more favorable to have a timer per namespace than using the algorithm you sent previously.
