

---

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

Posted by [den](#) on Thu, 18 Oct 2007 09:53:04 GMT

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

---

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).

The locking in mainstream is different with NETNS49.

I have proposed to add a dev\_base\_lock to protect namespace list for atomic context to be in sync with rtnl usage.

May be we should introduce an additional one. I also do not see other way for that place. May be you can copy my approach with generations for IPv6 code. I'll send a new version in a minute.

Regards,  
Den

---