Subject: Re: [PATCH] [NETNS49] support for per/namespace routing cache cleanup
Posted by dlunev on Thu, 18 Oct 2007 18:59:50 GMT
View Forum Message <> Reply to Message

Benjamin Thery wrote:
> 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 :)
>
> Why? :)

This is
- scalable: no long-long iteration over 1000 namespace
- easy in respect to a patch I just sent. We should not allow to drop
cache from all namespaces from one
- and this approach was good for Alexey Kuznetsov. We have talked about
this discussing OpenVZ implementation

> Then you'll have to find a way pass the target net to fib6_run_gc()
> (the timer handler). current->nsproxy->net_ns won't work :)
> One timer for all looked simpler to me.
>
> Can there be an impact on performance if we have several GC timers
> for the several netns running?

the amount of job is not greater. Isn't it?

Regards,
Den