

---

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

Posted by [ebiederm](#) on Fri, 19 Oct 2007 19:03:31 GMT

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

---

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

> Daniel Lezcano wrote:

>> 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.

>

> as do I. But the IPv4 and IPv6 situation is much different. IPv6 cache is  
> separated while IPv4 one is common.

>

> So, the situation is not uniform, at least :(

>

> By the way, we can move routing cache into the FIB tree like one for IPv4 :)

I'm trying to think through this.

- We have two cases garbage collection and flushing the cache.
- The routing cache is just a cache we loose nothing except performance if it is flushed. At least that is my memory.
- Especially for ipv4 the routing cache is a shared resource because of a shared hash table, therefore we want to do the garbage collection globally. Especially since we should be able to make better decisions that way.
- We appear to use a timer beneath flush to keep flushing operations from happening to frequently, and thus killing everyones performance.
- We also flush the timer when it is time to change secrets.

So as long as we are really only making the user space flush operation per namespace I don't have a real problem with it, as that timer should not run periodically and the whole reason for the existence of the timer is to ensure people don't kill other peoples performance by flushing to often.

I do still think things like the periodic secret rebuild and the garbage collection should continue to be global across the entire cache.

Eric

---