
Subject: Re: [PATCH net-next 0/9] selective (per/namespace) flush of rt_cache
Posted by [den](#) on Mon, 07 Jul 2008 10:39:41 GMT

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

On Mon, 2008-07-07 at 03:29 -0700, Eric W. Biederman wrote:

> David Miller <davem@davemloft.net> writes:

>

> > From: "Denis V. Lunev" <den@openvz.org>

> > Date: Fri, 04 Jul 2008 17:16:00 +0400

> >

> >> This series of patches implements selective rt cache flushing to make
> >> sure that in one namespace we'll not been able to affect the performance
> >> of other from the user space.

> >

> > Applied and pushed out to net-next-2.6, thanks.

> >

> > Although I wish patch 9 didn't have to be so ugly. :-/ Also, is it
> > really the right thing to do if another namespace's RT cache entries
> > are in fact chewing up all the slots in a hash chain? I think
> > the replacement garbage collection algorithm should be namespace
> > agnostic.

>

> My VM experience agrees. The requested flushes from inside the namespace
> and for namespace exit really should just flush the routes for that namespace.
> However for general route caching and expiry I don't see the point of making
> the code per namespace, and my gut feel is that is likely to reduce average
> case performance at the cost of better isolation between namespaces.

>

> Denis did I read the patches right and you are making all route cache flushes
> per namespace? Including the periodic expiry?

yes. The reason for the latter is that we can't pin namespace list in
the timer, so it is better to have a separate timer for each namespace.

The similar approach has been discussed for IPv6 flushing if my memory
does not fail me.

Though, generic rt cache garbage collecting is left untouched and from
my point of view it should be left as it is. At least we use this
approach in OpenVz.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
