
Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by [Carl-Daniel Hailfinge](#) on Sat, 23 Jun 2007 20:19:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 23.06.2007 19:19, Eric W. Biederman wrote:

> Patrick McHardy <kaber@trash.net> writes:

>

>> Eric W. Biederman wrote:

>

>>> Depending upon the data structure it will either be modified to hold
>>> a per entry network namespace pointer or it there will be a separate
>>> copy per network namespace. For large global data structures like
>>> the ipv4 routing cache hash table adding an additional pointer to the
>>> entries appears the more reasonable solution.

>>

>> So the routing cache is shared between all namespaces?

>

> Yes. Each namespaces has it's own view so semantically it's not
> shared. But the initial fan out of the hash table 2M or something
> isn't something we want to replicate on a per namespace basis even
> assuming the huge page allocations could happen.

>

> So we just tag the entries and add the network namespace as one more
> part of the key when doing hash table look ups.

Can one namespace DoS other namespaces' access to the routing cache?

Two scenarios come to mind:

* provoking hash collisions

* lock contention (sorry, haven't checked whether/how we do locking)

Regards,
Carl-Daniel

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
