
Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by [ebiederm](#) on Sat, 23 Jun 2007 19:08:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Patrick McHardy <kaber@trash.net> writes:

>>>> 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.
>
>
> I can wait for the patches, but I would be interested in how
> GC is performed and whether limits can be configured per
> namespace.

Currently I believe the gc code is unmodified in my patches.

Currently I have been focusing on the normal semantics and just making something work in a mergeable fashion.

Limits and the like are comparatively easy to add in after the rest is working so I haven't been focusing on that.

Eric

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