
Subject: Re: [patch 12/38][IPV6] ip6_fib - move the fib table to the network namespace

Posted by [Daniel Lezcano](#) on Wed, 05 Dec 2007 10:05:00 GMT

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

Brian Haley wrote:

> Hi Daniel,

>

> Daniel Lezcano wrote:

>> Move the global definition of the fib table to the network namespace
>> structure and make their access to the initial network namespace.

>

```
>> -void __init fib6_init(void)
>> +static int fib6_net_init(struct net *net)
>> {
>> -    fib6_node_kmem = kmem_cache_create("fib6_nodes",
>> -                                         sizeof(struct fib6_node),
>> -                                         0, SLAB_HWCACHE_ALIGN|SLAB_PANIC,
>> -                                         NULL);
>> +    int ret;
>>
>> -    fib_table_hash =
>> kzalloc(sizeof(*fib_table_hash)*FIB_TABLE_HASHSZ, GFP_KERNEL);
>> -    if (!fib_table_hash)
>> -        panic("IPV6: Failed to allocate fib_table_hash.\n");
>> -
>> -    fib6_main_tbl = kzalloc(sizeof(*fib6_main_tbl), GFP_KERNEL);
>> -    if (!fib6_main_tbl)
>> -        panic("IPV6: Failed to allocate fib6_main_tbl.\n");
>> -
>> -    fib6_main_tbl->tb6_id = RT6_TABLE_MAIN;
>> -    fib6_main_tbl->tb6_root.leaf = &ip6_null_entry;
>> -    fib6_main_tbl->tb6_root.fn_flags = RTN_ROOT | RTN_TL_ROOT |
>> RTN_RTINFO;
>> +    if (net != &init_net)
>> +        return -EPERM;
>> +
>> +    ret = -ENOMEM;
>> +    net->fib_table_hash =
>> kzalloc(sizeof(*net->fib_table_hash)*FIB_TABLE_HASHSZ,
>> +           GFP_KERNEL);
>> +    if (!net->fib_table_hash)
>> +        goto out;
```

> So originally, the fib_table_hash was global with no allocation

> necessary. Then in patch 11 you changed it to be dynamic and panic() on

> failure. Now it just gracefully returns. Do you really want to do that

> for the init_net namespace? Won't the next routine that tries to access

> this NULL pointer oops? Same for fib6_main_table, fib6_local_table,
> fib6_stats, etc. Or did I miss where the callers error-out on an ENOMEM?

Yes, you are right, this is not consistent and we should not panic
neither ignore the error. I should have modified the ip6_route_init and
returned an error in case of register_pernet_subsys failure. So the ipv6
initialization can fails safely.

I will post some modifications around that for net-2.6.25 for netdev@.

Thanks for catching that.

-- Daniel

Containers mailing list

Containers@lists.linux-foundation.org

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