
Subject: [patch 19/38][IPV6] ip6_fib - make the ip6 fib gc timer handle several network namespaces

Posted by [Daniel Lezcano](#) on Mon, 03 Dec 2007 16:16:55 GMT

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

Move the timer initialization at the network namespace creation and store the network namespace in the timer argument.

That enables multiple timers (one per network namespace) to do garbage collecting.

Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>

Signed-off-by: Benjamin Thery <benjamin.thery@bull.net>

net/ipv6/ip6_fib.c | 32 ++++++-----

1 file changed, 17 insertions(+), 15 deletions(-)

Index: linux-2.6-netns/net/ipv6/ip6_fib.c

=====

--- linux-2.6-netns.orig/net/ipv6/ip6_fib.c

+++ linux-2.6-netns/net/ipv6/ip6_fib.c

@@ -1471,12 +1471,23 @@ static void fib6_gc_timer_cb(unsigned lo
static int fib6_net_init(struct net *net)

{

int ret;

+ struct timer_list *timer;

ret = -ENOMEM;

+ timer = kzalloc(sizeof(*timer), GFP_KERNEL);

+ if (!timer)

+ goto out;

+

+ timer->function = fib6_gc_timer_cb;

+ timer->expires = 0;

+ timer->data = (unsigned long)net;

+ timer->base = &boot_tvec_bases;

+ net->ip6_fib_timer = timer;

+

net->fib_table_hash = kzalloc(sizeof(*net->fib_table_hash)*FIB_TABLE_HASHSZ,
GFP_KERNEL);

if (!net->fib_table_hash)

- goto out;

+ goto out_timer;

net->fib6_main_tbl = kzalloc(sizeof(*net->fib6_main_tbl), GFP_KERNEL);

if (!net->fib6_main_tbl)

@@ -1503,12 +1514,17 @@ static int fib6_net_init(struct net *net

```

out_fib6_main_tbl:
    kfree(net->fib_table_hash);
+out_timer:
+ kfree(timer);
out:
    return ret;
}

static void fib6_net_exit(struct net *net)
{
+ del_timer(net->ip6_fib_timer);
+ kfree(net->ip6_fib_timer);
+
#ifndef CONFIG_IPV6_MULTIPLE_TABLES
    kfree(net->fib6_local_tbl);
#endif
@@ -1523,31 +1539,17 @@ static struct pernet_operations fib6_net

void __init fib6_init(void)
{
- struct timer_list *timer;
-
    fib6_node_kmem = kmalloc(sizeof(struct fib6_node),
        sizeof(struct fib6_node),
        0, SLAB_HWCACHE_ALIGN|SLAB_PANIC,
        NULL);

- timer = kzalloc(sizeof(*timer), GFP_KERNEL);
- if (!timer)
-     panic("IPV6: failed to allocate the gc timer\n");
-
- timer->function = fib6_gc_timer_cb;
- timer->expires = 0;
- timer->data = (unsigned long)&init_net;
- timer->base = &boot_tvec_bases;
- init_net.ip6_fib_timer = timer;
-
    register_pernet_subsys(&fib6_net_ops);
    __ rtnl_register(PF_INET6, RTM_GETROUTE, NULL, inet6_dump_fib);
}

void fib6_gc_cleanup(void)
{
- del_timer(init_net.ip6_fib_timer);
- kfree(init_net.ip6_fib_timer);
    unregister_pernet_subsys(&fib6_net_ops);
    kmalloc_cache_destroy(fib6_node_kmem);
}

```

--

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