
Subject: [PATCH 3/3] [IPV6]: Fix refcounting for anycast dst entries.

Posted by [den](#) on Tue, 18 Mar 2008 14:35:25 GMT

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

Anycast DST entries allocated inside `ipv6_dev_ac_inc` are leaked when network device is stopped without removing IPv6 addresses from it. The bug has been observed in the reality on 2.6.18-rhel5 kernel.

In the above case `addrconf_ifdown` marks all entries as obsolete and `ip6_del_rt` called from `__ipv6_dev_ac_dec` returns `ENOENT`. The reference is not dropped.

The fix is simple. DST entry should not keep reference when stored in the FIB6 tree.

Signed-off-by: Denis V. Lunev <den@openvz.org>

net/ipv6/anycast.c | 9 ++-----

1 files changed, 2 insertions(+), 7 deletions(-)

diff --git a/net/ipv6/anycast.c b/net/ipv6/anycast.c

index 96868b9..7bc0469 100644

--- a/net/ipv6/anycast.c

+++ b/net/ipv6/anycast.c

@@ -334,9 +334,7 @@ int ipv6_dev_ac_inc(struct net_device *dev, struct in6_addr *addr)
 idev->ac_list = aca;
 write_unlock_bh(&idev->lock);

- dst_hold(&rt->u.dst);

- if (ip6_ins_rt(rt))

- dst_release(&rt->u.dst);

+ ip6_ins_rt(rt);

addrconf_join_solict(dev, &aca->aca_addr);

@@ -378,10 +376,7 @@ int __ipv6_dev_ac_dec(struct inet6_dev *idev, struct in6_addr *addr)
 addrconf_leave_solict(idev, &aca->aca_addr);

dst_hold(&aca->aca_rt->u.dst);

- if (ip6_del_rt(aca->aca_rt))

- dst_free(&aca->aca_rt->u.dst);

- else

- dst_release(&aca->aca_rt->u.dst);

+ ip6_del_rt(aca->aca_rt);

aca_put(aca);

return 0;

--

1.5.3.rc5

Containers mailing list

Containers@lists.linux-foundation.org

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