
Subject: [PATCH 6/6] [NETNS]: Lookup in FIB semantic hashes taking into account the namespace.

Posted by [den](#) on Thu, 31 Jan 2008 12:00:50 GMT

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

The namespace is not available in the fib_sync_down_addr, add it as a parameter.

Looking up a device by the pointer to it is OK. Looking up using a result from fib_trie/fib_hash table lookup is also safe. No need to fix that at all. So, just fix lookup by address and insertion to the hash table path.

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

```
include/net/ip_fib.h | 2 ++
net/ipv4/fib_frontend.c | 2 ++
net/ipv4/fib_semantics.c | 6 +++++-
3 files changed, 7 insertions(+), 3 deletions(-)
```

```
diff --git a/include/net/ip_fib.h b/include/net/ip_fib.h
index cb0df37..90d1175 100644
--- a/include/net/ip_fib.h
+++ b/include/net/ip_fib.h
@@ -220,7 +220,7 @@ extern void fib_select_default(struct net *net, const struct flowi *flp,
/* Exported by fib_semantics.c */
extern int ip_fib_check_default(__be32 gw, struct net_device *dev);
extern int fib_sync_down_dev(struct net_device *dev, int force);
-extern int fib_sync_down_addr(__be32 local);
+extern int fib_sync_down_addr(struct net *net, __be32 local);
extern int fib_sync_up(struct net_device *dev);
extern __be32 __fib_res_prefsrc(struct fib_result *res);
extern void fib_select_multipath(const struct flowi *flp, struct fib_result *res);
diff --git a/net/ipv4/fib_frontend.c b/net/ipv4/fib_frontend.c
index d69ffa2..86ff271 100644
--- a/net/ipv4/fib_frontend.c
+++ b/net/ipv4/fib_frontend.c
@@ -808,7 +808,7 @@ static void fib_del_ifaddr(struct in_ifaddr *ifa)
    First of all, we scan fib_info list searching
    for stray nexthop entries, then ignite fib_flush.
    */
-   if (fib_sync_down_addr(ifa->ifa_local))
+   if (fib_sync_down_addr(dev->nd_net, ifa->ifa_local))
       fib_flush(dev->nd_net);
}
}

diff --git a/net/ipv4/fib_semantics.c b/net/ipv4/fib_semantics.c
index 97cc494..a13c847 100644
--- a/net/ipv4/fib_semantics.c
```

```

+++ b/net/ipv4/fib_semantics.c
@@ -229,6 +229,8 @@ static struct fib_info *fib_find_info(const struct fib_info *nfi)
    head = &fib_info_hash[hash];

    hlist_for_each_entry(fi, node, head, fib_hash) {
+   if (fi->fib_net != nfi->fib_net)
+       continue;
    if (fi->fib_nhs != nfi->fib_nhs)
        continue;
    if (nfi->fib_protocol == fi->fib_protocol &&
@@ -1031,7 +1033,7 @@ nla_put_failure:
        referring to it.
    - device went down -> we must shutdown all nexthops going via it.
 */
-int fib_sync_down_addr(__be32 local)
+int fib_sync_down_addr(struct net *net, __be32 local)
{
    int ret = 0;
    unsigned int hash = fib_laddr_hashfn(local);
@@ -1043,6 +1045,8 @@ int fib_sync_down_addr(__be32 local)
    return 0;

    hlist_for_each_entry(fi, node, head, fib_lhash) {
+   if (fi->fib_net != net)
+       continue;
    if (fi->fib_prefsrc == local) {
        fi->fib_flags |= RTNH_F_DEAD;
        ret++;
    }
--
```

1.5.3.rc5
