
Subject: [PATCH 5/6] [NETNS]: Add a namespace mark to fib_info.
Posted by [den](#) on Thu, 31 Jan 2008 12:00:49 GMT

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

This is required to make fib_info lookups namespace aware. In the other case initial namespace devices are marked as dead in the local routing table during other namespace stop.

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

```
include/net/ip_fib.h    |  1 +
net/ipv4/fib_semantics.c |  8 ++++++-
2 files changed, 5 insertions(+), 4 deletions(-)
```

```
diff --git a/include/net/ip_fib.h b/include/net/ip_fib.h
index 1b2f008..cb0df37 100644
--- a/include/net/ip_fib.h
+++ b/include/net/ip_fib.h
@@ -69,6 +69,7 @@ struct fib_nh {
struct fib_info {
    struct hlist_node fib_hash;
    struct hlist_node fib_lhash;
+   struct net *fib_net;
    int fib_treeref;
    atomic_t fib_clntref;
    int fib_dead;
diff --git a/net/ipv4/fib_semantics.c b/net/ipv4/fib_semantics.c
index 5beff2e..97cc494 100644
--- a/net/ipv4/fib_semantics.c
+++ b/net/ipv4/fib_semantics.c
@@ -687,6 +687,7 @@ struct fib_info *fib_create_info(struct fib_config *cfg)
    struct fib_info *fi = NULL;
    struct fib_info *ofi;
    int nhs = 1;
+   struct net *net = cfg->fc_nlinfo.nl_net;

/* Fast check to catch the most weird cases */
if (fib_props[cfg->fc_type].scope > cfg->fc_scope)
@@ -727,6 +728,7 @@ struct fib_info *fib_create_info(struct fib_config *cfg)
    goto failure;
    fib_info_cnt++;

+ fi->fib_net = net;
    fi->fib_protocol = cfg->fc_protocol;
    fi->fib_flags = cfg->fc_flags;
    fi->fib_priority = cfg->fc_priority;
@@ -798,8 +800,7 @@ struct fib_info *fib_create_info(struct fib_config *cfg)
    if (nhs != 1 || nh->nh_gw)
```

```
    goto err_inval;
    nh->nh_scope = RT_SCOPE_NOWHERE;
- nh->nh_dev = dev_get_by_index(cfg->fc_nlinfo.nl_net,
-     fi->fib_nh->nh_oif);
+ nh->nh_dev = dev_get_by_index(net, fi->fib_nh->nh_oif);
    err = -ENODEV;
    if (nh->nh_dev == NULL)
        goto failure;
@@ @ -813,8 +814,7 @@ struct fib_info *fib_create_info(struct fib_config *cfg)
if (fi->fib_prefsrc) {
    if (cfg->fc_type != RTN_LOCAL || !cfg->fc_dst ||
        fi->fib_prefsrc != cfg->fc_dst)
-     if (inet_addr_type(cfg->fc_nlinfo.nl_net,
-         fi->fib_prefsrc) != RTN_LOCAL)
+     if (inet_addr_type(net, fi->fib_prefsrc) != RTN_LOCAL)
        goto err_inval;
}
```

--

1.5.3.rc5
