
Subject: [PATCH 3/12 net-2.6.25] [NETNS]: Add netns parameter to fib_select_default.

Posted by [den](#) on Tue, 22 Jan 2008 15:59:20 GMT

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

Currently fib_select_default calls fib_get_table() with the init_net. Prepare it to provide a correct namespace to lookup default route.

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

```
include/net/ip_fib.h | 3 ++
net/ipv4/fib_frontend.c | 5 +++
net/ipv4/route.c      | 2 ++
3 files changed, 6 insertions(+), 4 deletions(-)
```

```
diff --git a/include/net/ip_fib.h b/include/net/ip_fib.h
index 39f944a..9daa60b 100644
--- a/include/net/ip_fib.h
+++ b/include/net/ip_fib.h
@@ -213,7 +213,8 @@ extern const struct nla_policy rtm_ipv4_policy[]; 
extern void ip_fib_init(void);
extern int fib_validate_source(__be32 src, __be32 dst, u8 tos, int oif,
    struct net_device *dev, __be32 *spec_dst, u32 *itag);
-extern void fib_select_default(const struct flowi *flp, struct fib_result *res);
+extern void fib_select_default(struct net *net, const struct flowi *flp,
+    struct fib_result *res);

/* Exported by fib_semantics.c */
extern int ip_fib_check_default(__be32 gw, struct net_device *dev);
diff --git a/net/ipv4/fib_frontend.c b/net/ipv4/fib_frontend.c
index 13bf01d..7e3e732 100644
--- a/net/ipv4/fib_frontend.c
+++ b/net/ipv4/fib_frontend.c
@@ -116,7 +116,8 @@ struct fib_table *fib_get_table(struct net *net, u32 id)
}
#endif /* CONFIG_IP_MULTIPLE_TABLES */

-void fib_select_default(const struct flowi *flp, struct fib_result *res)
+void fib_select_default(struct net *net,
+    const struct flowi *flp, struct fib_result *res)
{
    struct fib_table *tb;
    int table = RT_TABLE_MAIN;
@@ -125,7 +126,7 @@ void fib_select_default(const struct flowi *flp, struct fib_result *res)
    return;
    table = res->r->table;
#endif
-    tb = fib_get_table(&init_net, table);
```

```
+ tb = fib_get_table(net, table);
if (FIB_RES_GW(*res) && FIB_RES_NH(*res).nh_scope == RT_SCOPE_LINK)
    tb->tb_select_default(tb, flp, res);
}
diff --git a/net/ipv4/route.c b/net/ipv4/route.c
index 27e0f81..4313255 100644
--- a/net/ipv4/route.c
+++ b/net/ipv4/route.c
@@@ -2419,7 +2419,7 @@ static int ip_route_output_slow(struct rtable **rp, const struct flowi
*oldflp)
else
#endif
if (!res.prefixlen && res.type == RTN_UNICAST && !fl.oif)
- fib_select_default(&fl, &res);
+ fib_select_default(&init_net, &fl, &res);

if (!fl.fl4_src)
fl.fl4_src = FIB_RES_PREFSRC(res);
--
```

1.5.3.rc5
