Subject: [PATCH 4/5] netns netfilter: per-netns arp_tables FILTER
Posted by Alexey Dobriyan on Thu, 24 Jan 2008 12:29:30 GMT
View Forum Message <> Reply to Message

Signed-off-by: Alexey Dobriyan <adobriyan@sw.ru>
---

 include/net/netns/ipv4.h | 1
 net/ipv4/netfilter/arptable_filter.c | 38 ++++++++++++++++++++++++++----------
 2 files changed, 29 insertions(+), 10 deletions(-)

```
--- a/include/net/netns/ipv4.h
+++ b/include/net/netns/ipv4.h
@@ -32,6 +32,7 @@ struct netns_ipv4 {
  struct xt_table  *iptable_filter;
  struct xt_table  *iptable_mangle;
  struct xt_table  *iptable_raw;
+ struct xt_table  *arptable_filter;
 #endif
 };
 #endif
--- a/net/ipv4/netfilter/arptable_filter.c
+++ b/net/ipv4/netfilter/arptable_filter.c
@@ -20,7 +20,7 @@ static struct
  struct arpt_replace repl;
  struct arpt_standard entries[3];
  struct arpt_error term;
-} initial_table __initdata = {
+} initial_table __net_initdata = {
  .repl = {
   .name = "filter",
   .valid_hooks = FILTER_VALID_HOOKS,
@@ -45,7 +45,7 @@ static struct
  .term = ARPT_ERROR_INIT,
 };

-static struct arpt_table __packet_filter = {
+static struct arpt_table packet_filter = {
  .name  = "filter",
  .valid_hooks = FILTER_VALID_HOOKS,
  .lock  = RW_LOCK_UNLOCKED,
@@ -53,7 +53,6 @@ static struct arpt_table __packet_filter = {
  .me  = THIS_MODULE,
  .af  = NF_ARP,
 };
-static struct arpt_table *packet_filter;

 /* The work comes in here from netfilter.c */
```

```
 static unsigned int arpt_hook(unsigned int hook,
@@ -62,7 +61,7 @@ static unsigned int arpt_hook(unsigned int hook,
         const struct net_device *out,
         int (*okfn)(struct sk_buff *))
 {
- return arpt_do_table(skb, hook, in, out, packet_filter);
+ return arpt_do_table(skb, hook, in, out, init_net.ipv4.arptable_filter);
 }

 static struct nf_hook_ops arpt_ops[] __read_mostly = {
@@ -86,14 +85,33 @@ static struct nf_hook_ops arpt_ops[] __read_mostly = {
  },
 };

+static int __net_init arptable_filter_net_init(struct net *net)
+{
+ /* Register table */
+ net->ipv4.arptable_filter =
+  arpt_register_table(net, &packet_filter, &initial_table.repl);
+ if (IS_ERR(net->ipv4.arptable_filter))
+  return PTR_ERR(net->ipv4.arptable_filter);
+ return 0;
+}
+
+static void __net_exit arptable_filter_net_exit(struct net *net)
+{
+ arpt_unregister_table(net->ipv4.arptable_filter);
+}
+
+static struct pernet_operations arptable_filter_net_ops = {
+ .init = arptable_filter_net_init,
+ .exit = arptable_filter_net_exit,
+};
+
 static int __init arptable_filter_init(void)
 {
  int ret;

- /* Register table */
- packet_filter = arpt_register_table(&init_net, &__packet_filter, &initial_table.repl);
- if (IS_ERR(packet_filter))
-  return PTR_ERR(packet_filter);
+ ret = register_pernet_subsys(&arptable_filter_net_ops);
+ if (ret < 0)
+  return ret;

  ret = nf_register_hooks(arpt_ops, ARRAY_SIZE(arpt_ops));
  if (ret < 0)
```

```
@@ -101,14 +119,14 @@ static int __init arptable_filter_init(void)
  return ret;

 cleanup_table:
- arpt_unregister_table(packet_filter);
+ unregister_pernet_subsys(&arptable_filter_net_ops);
  return ret;
 }

 static void __exit arptable_filter_fini(void)
 {
  nf_unregister_hooks(arpt_ops, ARRAY_SIZE(arpt_ops));
- arpt_unregister_table(packet_filter);
+ unregister_pernet_subsys(&arptable_filter_net_ops);
 }

 module_init(arptable_filter_init);
```