
Subject: [PATCH net-2.6.25 10/10][NETNS][FRAGS]: Make the pernet subsystem for fragments.

Posted by Pavel Emelianov on Tue, 22 Jan 2008 14:10:51 GMT

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

On namespace start we mainly prepare the ctl variables.

When the namespace is stopped we have to kill all the fragments that point to this namespace.

The inet_frags_exit_net() handles it.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

```
include/net/inet_frag.h |  1 +
net/ipv4/inet_fragment.c |  7 ++++++
net/ipv4/ip_fragment.c  | 13 ++++++++++++
net/ipv6/reassembly.c   | 14 ++++++++++++
4 files changed, 33 insertions(+), 2 deletions(-)
```

```
diff --git a/include/net/inet_frag.h b/include/net/inet_frag.h
index 3695ff4..7374251 100644
--- a/include/net/inet_frag.h
+++ b/include/net/inet_frag.h
@@ -54,6 +54,7 @@ void inet_frags_init(struct inet_frags *);
void inet_frags_fini(struct inet_frags *);

void inet_frags_init_net(struct netns_frags *nf);
+void inet_frags_exit_net(struct netns_frags *nf, struct inet_frags *f);

void inet_frag_kill(struct inet_frag_queue *q, struct inet_frags *f);
void inet_frag_destroy(struct inet_frag_queue *q,
diff --git a/net/ipv4/inet_fragment.c b/net/ipv4/inet_fragment.c
index f1b95e1..724d69a 100644
--- a/net/ipv4/inet_fragment.c
+++ b/net/ipv4/inet_fragment.c
@@ -83,6 +83,13 @@ void inet_frags_fini(struct inet_frags *f)
}
EXPORT_SYMBOL(inet_frags_fini);

+void inet_frags_exit_net(struct netns_frags *nf, struct inet_frags *f)
+{
+ nf->low_thresh = 0;
+ inet_frag_evictor(nf, f);
+}
+EXPORT_SYMBOL(inet_frags_exit_net);
+
static inline void fq_unlink(struct inet_frag_queue *fq, struct inet_frags *f)
```

```

{
 write_lock(&f->lock);
diff --git a/net/ipv4/ip_fragment.c b/net/ipv4/ip_fragment.c
index 29b4b09..a2e92f9 100644
--- a/net/ipv4/ip_fragment.c
+++ b/net/ipv4/ip_fragment.c
@@ -713,9 +713,20 @@ static int ipv4_frags_init_net(struct net *net)
 return ip4_frags_ctl_register(net);
}

+static void ipv4_frags_exit_net(struct net *net)
+{
+ ip4_frags_ctl_unregister(net);
+ inet_frags_exit_net(&net->ipv4.frags, &ip4_frags);
+}
+
+static struct pernet_operations ip4_frags_ops = {
+ .init = ipv4_frags_init_net,
+ .exit = ipv4_frags_exit_net,
+};
+
void __init ipfrag_init(void)
{
- ipv4_frags_init_net(&init_net);
+ register_pernet_subsys(&ip4_frags_ops);
 ip4_frags.hashfn = ip4_hashfn;
 ip4_frags.constructor = ip4_frag_init;
 ip4_frags.destructor = ip4_frag_free;
diff --git a/net/ipv6/reassembly.c b/net/ipv6/reassembly.c
index 0c4bc46..f936d04 100644
--- a/net/ipv6/reassembly.c
+++ b/net/ipv6/reassembly.c
@@ -728,6 +728,17 @@ static int ipv6_frags_init_net(struct net *net)
 return ip6_frags_sysctl_register(net);
}

+static void ipv6_frags_exit_net(struct net *net)
+{
+ ip6_frags_sysctl_unregister(net);
+ inet_frags_exit_net(&net->ipv6.frags, &ip6_frags);
+}
+
+static struct pernet_operations ip6_frags_ops = {
+ .init = ipv6_frags_init_net,
+ .exit = ipv6_frags_exit_net,
+};
+
int __init ipv6_frag_init(void)

```

```
{  
    int ret;  
@@ -736,7 +747,7 @@ int __init ipv6_frag_init(void)  
if (ret)  
    goto out;  
  
- ipv6 frags_init_net(&init_net);  
+ register_pernet_subsys(&ip6 frags_ops);  
  
ip6 frags.hashfn = ip6 hashfn;  
ip6 frags.constructor = ip6 frag_init;  
@@ -754,5 +765,6 @@ out:  
void ipv6_frag_exit(void)  
{  
    inet frags_fini(&ip6 frags);  
+ unregister_pernet_subsys(&ip6 frags_ops);  
    inet6 del_protocol(&frag_protocol, IPPROTO_FRAGMENT);  
}  
--
```

1.5.3.4
