
Subject: [PATCH net-2.6.25 3/10][NETNS][FRAGS]: Make the nqueues counter per-namespace.

Posted by Pavel Emelianov on Tue, 22 Jan 2008 13:58:13 GMT

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

This is simple - just move the variable from struct `inet_frag.h` to struct `netns_frag.h` and adjust the usage appropriately.

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

```
include/net/inet_frag.h      |  4 +---  
include/net/ip.h            |  2 +-  
include/net/ipv6.h          |  2 +-  
net/ipv4/inet_fragment.c   | 11 ++++++----  
net/ipv4/ip_fragment.c     |  6 +----  
net/ipv4/proc.c            |  2 +-  
net/ipv6/netfilter/nf_conntrack_reasm.c |  1 +  
net/ipv6/proc.c            |  2 +-  
net/ipv6/reassembly.c      |  6 +----  
9 files changed, 24 insertions(+), 12 deletions(-)
```

```
diff --git a/include/net/inet_frag.h b/include/net/inet_frag.h  
index 8ab6df6..d36f3a6 100644  
--- a/include/net/inet_frag.h  
+++ b/include/net/inet_frag.h  
@@ -2,6 +2,7 @@  
#define __NET_FRAG_H__
```

```
struct netns_frag {  
+ int nqueues;  
};  
  
struct inet_frag_queue {  
@@ -36,7 +37,6 @@ struct inet_frag {  
    struct hlist_head hash[INETFRAGS_HASHSZ];  
    rwlock_t lock;  
    u32 rnd;  
- int nqueues;  
    int qsize;  
    atomic_t mem;  
    struct timer_list secret_timer;  
@@ -55,6 +55,8 @@ struct inet frags {  
    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_frag_kill(struct inet_frag_queue *q, struct inet_frags *f);
void inet_frag_destroy(struct inet_frag_queue *q,
                      struct inet_frags *f, int *work);
diff --git a/include/net/ip.h b/include/net/ip.h
index ff14fc8..9ea1bc5 100644
--- a/include/net/ip.h
+++ b/include/net/ip.h
@@ -330,7 +330,7 @@ enum ip_defrag_users

int ip_defrag(struct sk_buff *skb, u32 user);
int ip_frag_mem(void);
-int ip_frag_nqueues(void);
+int ip_frag_nqueues(struct net *net);

/*
 * Functions provided by ip_forward.c
diff --git a/include/net/ipv6.h b/include/net/ipv6.h
index 87ca1bf..da1c089 100644
--- a/include/net/ipv6.h
+++ b/include/net/ipv6.h
@@ -245,7 +245,7 @@ struct ipv6_txoptions *ipv6_fixup_options(struct ipv6_txoptions
*opt_space,
extern int ipv6_opt_accepted(struct sock *sk, struct sk_buff *skb);

-int ip6_frag_nqueues(void);
+int ip6_frag_nqueues(struct net *net);
int ip6_frag_mem(void);

#define IPV6_FRAG_TIMEOUT (60*HZ) /* 60 seconds */
diff --git a/net/ipv4/inet_fragment.c b/net/ipv4/inet_fragment.c
index 158c5f6..4fec0b9 100644
--- a/net/ipv4/inet_fragment.c
+++ b/net/ipv4/inet_fragment.c
@@ -63,7 +63,6 @@ void inet_frags_init(struct inet_frags *f)
    f->rnd = (u32) ((num_physpages ^ (num_physpages>>7)) ^
                     (jiffies ^ (jiffies >> 6)));
}

-f->nqueues = 0;
atomic_set(&f->mem, 0);

setup_timer(&f->secret_timer, inet_frag_secret_rebuild,
@@ -73,6 +72,12 @@ void inet_frags_init(struct inet_frags *f)
}
EXPORT_SYMBOL(inet_frags_init);

+void inet_frags_init_net(struct netns_frags *nf)
+{

```

```

+ nf->nqueues = 0;
+}
+EXPORT_SYMBOL(inet_frags_init_net);
+
void inet_frags_fini(struct inet_frags *f)
{
    del_timer(&f->secret_timer);
@@ -84,7 +89,7 @@ static inline void fq_unlink(struct inet_frag_queue *fq, struct inet_frags *f)
    write_lock(&f->lock);
    hlist_del(&fq->list);
    list_del(&fq->lru_list);
- f->nqueues--;
+ fq->net->nqueues--;
    write_unlock(&f->lock);
}

@@ -206,7 +211,7 @@ static struct inet_frag_queue *inet_frag_intern(struct netns_frags *nf,
    atomic_inc(&qp->refcnt);
    hlist_add_head(&qp->list, &f->hash[hash]);
    list_add_tail(&qp->lru_list, &f->lru_list);
- f->nqueues++;
+ nf->nqueues++;
    write_unlock(&f->lock);
    return qp;
}
diff --git a/net/ipv4/ip_fragment.c b/net/ipv4/ip_fragment.c
index 56211ef..cd8c830 100644
--- a/net/ipv4/ip_fragment.c
+++ b/net/ipv4/ip_fragment.c
@@ -95,9 +95,9 @@ static struct inet_frags_ctl ip4_frags_ctl __read_mostly = {

static struct inet_frags ip4_frags;

-int ip_frag_nqueues(void)
+int ip_frag_nqueues(struct net *net)
{
- return ip4_frags.nqueues;
+ return net->ipv4.frags.nqueues;
}

int ip_frag_mem(void)
@@ -675,6 +675,8 @@ static inline int ip4_frags_ctl_register(struct net *net)

static int ipv4_frags_init_net(struct net *net)
{
+ inet_frags_init_net(&net->ipv4.frags);
+
    return ip4_frags_ctl_register(net);

```

```

}

diff --git a/net/ipv4/proc.c b/net/ipv4/proc.c
index cb3787f..bae3280 100644
--- a/net/ipv4/proc.c
+++ b/net/ipv4/proc.c
@@ -62,7 +62,7 @@ static int sockstat_seq_show(struct seq_file *seq, void *v)
    seq_printf(seq, "UDPLITE: inuse %d\n", sock_prot_inuse_get(&udlite_prot));
    seq_printf(seq, "RAW: inuse %d\n", sock_prot_inuse_get(&raw_prot));
    seq_printf(seq, "FRAG: inuse %d memory %d\n",
-   ip_frag_nqueues(), ip_frag_mem());
+   ip_frag_nqueues(&init_net), ip_frag_mem());
    return 0;
}

diff --git a/net/ipv6/netfilter/nf_conntrack_reasm.c b/net/ipv6/netfilter/nf_conntrack_reasm.c
index 18accd4..0b9d009 100644
--- a/net/ipv6/netfilter/nf_conntrack_reasm.c
+++ b/net/ipv6/netfilter/nf_conntrack_reasm.c
@@ -712,6 +712,7 @@ int nf_ct_frag6_init(void)
    nf frags.qsize = sizeof(struct nf_ct_frag6_queue);
    nf frags.match = ip6_frag_match;
    nf frags.frag_expire = nf_ct_frag6_expire;
+ inet frags_init_net(&nf_init frags);
    inet frags_init(&nf frags);

    return 0;
diff --git a/net/ipv6/proc.c b/net/ipv6/proc.c
index 6b0314e..0b55785 100644
--- a/net/ipv6/proc.c
+++ b/net/ipv6/proc.c
@@ -44,7 +44,7 @@ static int sockstat6_seq_show(struct seq_file *seq, void *v)
    seq_printf(seq, "RAW6: inuse %d\n",
               sock_prot_inuse_get(&rawv6_prot));
    seq_printf(seq, "FRAG6: inuse %d memory %d\n",
-   ip6_frag_nqueues(), ip6_frag_mem());
+   ip6_frag_nqueues(&init_net), ip6_frag_mem());
    return 0;
}

diff --git a/net/ipv6/reassembly.c b/net/ipv6/reassembly.c
index ab2d53b..77a8740 100644
--- a/net/ipv6/reassembly.c
+++ b/net/ipv6/reassembly.c
@@ -84,9 +84,9 @@ struct frag_queue

static struct inet frags ip6 frags;

```

```
-int ip6_frag_nqueues(void)
+int ip6_frag_nqueues(struct net *net)
{
- return ip6 frags.nqueues;
+ return net->ipv6 frags.nqueues;
}

int ip6_frag_mem(void)
@@ -690,6 +690,8 @@ static int ipv6 frags_init_net(struct net *net)
net->ipv6 sysctl frags.timeout = IPV6_FRAG_TIMEOUT;
net->ipv6 sysctl frags.secret_interval = 10 * 60 * HZ;

+ inet frags_init_net(&net->ipv6 frags);
+
return ip6 frags sysctl_register(net);
}

--
```

1.5.3.4
