
Subject: [PATCH 3/3] netns netfilter: create per-netns /proc/net/*_tables_*

Posted by [Alexey Dobriyan](#) on Fri, 25 Jan 2008 16:45:29 GMT

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

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

```
include/linux/netfilter/x_tables.h | 4 +---
net/ipv4/netfilter/arp_tables.c    | 21 ++++++++
net/ipv4/netfilter/ip_tables.c     | 21 ++++++++
net/ipv6/netfilter/ip6_tables.c    | 22 ++++++++
net/netfilter/x_tables.c           | 20 ++++++-----
5 files changed, 67 insertions(+), 21 deletions(-)
```

--- a/include/linux/netfilter/x_tables.h

+++ b/include/linux/netfilter/x_tables.h

```
@ @ -357,8 +357,8 @ @ extern struct xt_table *xt_find_table_lock(struct net *net, int af,
    const char *name);
extern void xt_table_unlock(struct xt_table *t);
```

```
-extern int xt_proto_init(int af);
```

```
-extern void xt_proto_fini(int af);
```

```
+extern int xt_proto_init(struct net *net, int af);
```

```
+extern void xt_proto_fini(struct net *net, int af);
```

```
extern struct xt_table_info *xt_alloc_table_info(unsigned int size);
```

```
extern void xt_free_table_info(struct xt_table_info *info);
```

--- a/net/ipv4/netfilter/arp_tables.c

+++ b/net/ipv4/netfilter/arp_tables.c

```
@ @ -1822,11 +1822,26 @ @ static struct nf_sockopt_ops arpt_sockopts = {
    .owner = THIS_MODULE,
};
```

```
+static int __net_init arp_tables_net_init(struct net *net)
```

```
+{
```

```
+ return xt_proto_init(net, NF_ARP);
```

```
+}
```

```
+
```

```
+static void __net_exit arp_tables_net_exit(struct net *net)
```

```
+{
```

```
+ xt_proto_fini(net, NF_ARP);
```

```
+}
```

```
+
```

```
+static struct pernet_operations arp_tables_net_ops = {
```

```
+ .init = arp_tables_net_init,
```

```
+ .exit = arp_tables_net_exit,
```

```
+};
```

```
+
```

```

static int __init arp_tables_init(void)
{
    int ret;

- ret = xt_proto_init(NF_ARP);
+ ret = register_pernet_subsys(&arp_tables_net_ops);
    if (ret < 0)
        goto err1;

@@ -1851,7 +1866,7 @@ err4:
err3:
    xt_unregister_target(&arpt_standard_target);
err2:
- xt_proto_fini(NF_ARP);
+ unregister_pernet_subsys(&arp_tables_net_ops);
err1:
    return ret;
}

@@ -1861,7 +1876,7 @@ static void __exit arp_tables_fini(void)
    nf_unregister_sockopt(&arpt_sockopts);
    xt_unregister_target(&arpt_error_target);
    xt_unregister_target(&arpt_standard_target);
- xt_proto_fini(NF_ARP);
+ unregister_pernet_subsys(&arp_tables_net_ops);
}

EXPORT_SYMBOL(arpt_register_table);
--- a/net/ipv4/netfilter/ip_tables.c
+++ b/net/ipv4/netfilter/ip_tables.c
@@ -2211,11 +2211,26 @@ static struct xt_match icmp_matchstruct __read_mostly = {
    .family = AF_INET,
};

+static int __net_init ip_tables_net_init(struct net *net)
+{
+ return xt_proto_init(net, AF_INET);
+}
+
+static void __net_exit ip_tables_net_exit(struct net *net)
+{
+ xt_proto_fini(net, AF_INET);
+}
+
+static struct pernet_operations ip_tables_net_ops = {
+ .init = ip_tables_net_init,
+ .exit = ip_tables_net_exit,
+};
+

```

```

static int __init ip_tables_init(void)
{
    int ret;

- ret = xt_proto_init(AF_INET);
+ ret = register_pernet_subsys(&ip_tables_net_ops);
    if (ret < 0)
        goto err1;

@@ -2245,7 +2260,7 @@ err4:
err3:
    xt_unregister_target(&ipt_standard_target);
err2:
- xt_proto_fini(AF_INET);
+ unregister_pernet_subsys(&ip_tables_net_ops);
err1:
    return ret;
}
@@ -2258,7 +2273,7 @@ static void __exit ip_tables_fini(void)
    xt_unregister_target(&ipt_error_target);
    xt_unregister_target(&ipt_standard_target);

- xt_proto_fini(AF_INET);
+ unregister_pernet_subsys(&ip_tables_net_ops);
}

EXPORT_SYMBOL(ipt_register_table);
--- a/net/ipv6/netfilter/ip6_tables.c
+++ b/net/ipv6/netfilter/ip6_tables.c
@@ -2235,11 +2235,26 @@ static struct xt_match icmp6_matchstruct __read_mostly = {
    .family = AF_INET6,
};

+static int __net_init ip6_tables_net_init(struct net *net)
+{
+ return xt_proto_init(net, AF_INET6);
+}
+
+static void __net_exit ip6_tables_net_exit(struct net *net)
+{
+ xt_proto_fini(net, AF_INET6);
+}
+
+static struct pernet_operations ip6_tables_net_ops = {
+ .init = ip6_tables_net_init,
+ .exit = ip6_tables_net_exit,
+};
+

```

```

static int __init ip6_tables_init(void)
{
    int ret;

- ret = xt_proto_init(AF_INET6);
+ ret = register_pernet_subsys(&ip6_tables_net_ops);
    if (ret < 0)
        goto err1;

@@ -2269,7 +2284,7 @@ err4:
err3:
    xt_unregister_target(&ip6t_standard_target);
err2:
- xt_proto_fini(AF_INET6);
+ unregister_pernet_subsys(&ip6_tables_net_ops);
err1:
    return ret;
}

@@ -2281,7 +2296,8 @@ static void __exit ip6_tables_fini(void)
    xt_unregister_match(&icmp6_matchstruct);
    xt_unregister_target(&ip6t_error_target);
    xt_unregister_target(&ip6t_standard_target);
- xt_proto_fini(AF_INET6);
+
+ unregister_pernet_subsys(&ip6_tables_net_ops);
}

/*
--- a/net/netfilter/x_tables.c
+++ b/net/netfilter/x_tables.c
@@ -928,7 +928,7 @@ static const struct file_operations xt_target_ops = {

#endif /* CONFIG_PROC_FS */

-int xt_proto_init(int af)
+int xt_proto_init(struct net *net, int af)
{
#ifdef CONFIG_PROC_FS
    char buf[XT_FUNCTION_MAXNAMELEN];
@@ -942,7 +942,7 @@ int xt_proto_init(int af)
#ifdef CONFIG_PROC_FS
    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_TABLES, sizeof(buf));
- proc = proc_net_fops_create(&init_net, buf, 0440, &xt_table_ops);
+ proc = proc_net_fops_create(net, buf, 0440, &xt_table_ops);
    if (!proc)
        goto out;
    proc->data = (void *) (unsigned long) af;

```

```

@@ -950,14 +950,14 @@ int xt_proto_init(int af)

    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_MATCHES, sizeof(buf));
- proc = proc_net_fops_create(&init_net, buf, 0440, &xt_match_ops);
+ proc = proc_net_fops_create(net, buf, 0440, &xt_match_ops);
    if (!proc)
        goto out_remove_tables;
    proc->data = (void *) (unsigned long)af;

    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_TARGETS, sizeof(buf));
- proc = proc_net_fops_create(&init_net, buf, 0440, &xt_target_ops);
+ proc = proc_net_fops_create(net, buf, 0440, &xt_target_ops);
    if (!proc)
        goto out_remove_matches;
    proc->data = (void *) (unsigned long)af;
@@ -969,34 +969,34 @@ int xt_proto_init(int af)
out_remove_matches:
    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_MATCHES, sizeof(buf));
- proc_net_remove(&init_net, buf);
+ proc_net_remove(net, buf);

out_remove_tables:
    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_TABLES, sizeof(buf));
- proc_net_remove(&init_net, buf);
+ proc_net_remove(net, buf);
out:
    return -1;
#endif
}
EXPORT_SYMBOL_GPL(xt_proto_init);

-void xt_proto_fini(int af)
+void xt_proto_fini(struct net *net, int af)
{
#ifdef CONFIG_PROC_FS
    char buf[XT_FUNCTION_MAXNAMELEN];

    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_TABLES, sizeof(buf));
- proc_net_remove(&init_net, buf);
+ proc_net_remove(net, buf);

    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_TARGETS, sizeof(buf));

```

```
- proc_net_remove(&init_net, buf);
+ proc_net_remove(net, buf);

    strcpy(buf, xt_prefix[af], sizeof(buf));
    strcat(buf, FORMAT_MATCHES, sizeof(buf));
- proc_net_remove(&init_net, buf);
+ proc_net_remove(net, buf);
#endif /*CONFIG_PROC_FS*/
}
EXPORT_SYMBOL_GPL(xt_proto_fini);
```

Subject: Re: [PATCH 3/3] netns netfilter: create per-netns /proc/net/*_tables_*
Posted by [Patrick McHardy](#) on Wed, 30 Jan 2008 15:17:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

Alexey Dobriyan wrote:

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

> ---

>

> include/linux/netfilter/x_tables.h | 4 +---

> net/ipv4/netfilter/arp_tables.c | 21 ++++++++-----

> net/ipv4/netfilter/ip_tables.c | 21 ++++++++-----

> net/ipv6/netfilter/ip6_tables.c | 22 ++++++++-----

> net/netfilter/x_tables.c | 20 ++++++++-----

> 5 files changed, 67 insertions(+), 21 deletions(-)

Also applied, thanks.
