
Subject: [PATCH net-2.6.25 2/4][NETNS][RAW]: Make /proc/net/raw(6) show per-namespace socket list.

Posted by [Pavel Emelianov](#) on Mon, 14 Jan 2008 13:08:42 GMT

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

Pull the struct net pointer up to the showing functions to filter the sockets depending on their namespaces.

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

```
---
include/net/raw.h | 3 +-
net/ipv4/raw.c   | 20 ++++++-----
net/ipv6/raw.c   | 4 +-
3 files changed, 16 insertions(+), 11 deletions(-)
```

```
diff --git a/include/net/raw.h b/include/net/raw.h
index 4d1aba0..cca81d8 100644
```

```
--- a/include/net/raw.h
```

```
+++ b/include/net/raw.h
```

```
@@ -39,6 +39,7 @@ extern int raw_proc_init(void);
extern void raw_proc_exit(void);
```

```
struct raw_iter_state {
```

```
+ struct seq_net_private p;
```

```
int bucket;
```

```
unsigned short family;
```

```
struct raw_hashinfo *h;
```

```
@@ -48,7 +49,7 @@ struct raw_iter_state {
```

```
void *raw_seq_start(struct seq_file *seq, loff_t *pos);
```

```
void *raw_seq_next(struct seq_file *seq, void *v, loff_t *pos);
```

```
void raw_seq_stop(struct seq_file *seq, void *v);
```

```
-int raw_seq_open(struct file *file, struct raw_hashinfo *h,
```

```
+int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h,
    unsigned short family);
```

```
#endif
```

```
diff --git a/net/ipv4/raw.c b/net/ipv4/raw.c
```

```
index a490a9d..4e95372 100644
```

```
--- a/net/ipv4/raw.c
```

```
+++ b/net/ipv4/raw.c
```

```
@@ -860,7 +860,8 @@ static struct sock *raw_get_first(struct seq_file *seq)
    struct hlist_node *node;
```

```
    sk_for_each(sk, node, &state->h->ht[state->bucket])
```

```
- if (sk->sk_family == state->family)
```

```
+ if (sk->sk_net == state->p.net &&
```

```
+ sk->sk_family == state->family)
```

```

    goto found;
}
sk = NULL;
@@ -876,7 +877,8 @@ static struct sock *raw_get_next(struct seq_file *seq, struct sock *sk)
    sk = sk_next(sk);
try_again:
;
- } while (sk && sk->sk_family != state->family);
+ } while (sk && sk->sk_net != state->p.net &&
+ sk->sk_family != state->family);

if (!sk && ++state->bucket < RAW_HTABLE_SIZE) {
    sk = sk_head(&state->h->ht[state->bucket]);
@@ -970,16 +972,18 @@ static const struct seq_operations raw_seq_ops = {
    .show = raw_seq_show,
};

-int raw_seq_open(struct file *file, struct raw_hashinfo *h,
+int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h,
    unsigned short family)
{
+ int err;
    struct raw_iter_state *i;

- i = __seq_open_private(file, &raw_seq_ops,
+ err = seq_open_net(ino, file, &raw_seq_ops,
    sizeof(struct raw_iter_state));
- if (i == NULL)
- return -ENOMEM;
+ if (err < 0)
+ return err;

+ i = raw_seq_private((struct seq_file *)file->private_data);
    i->h = h;
    i->family = family;
    return 0;
@@ -988,7 +992,7 @@ EXPORT_SYMBOL_GPL(raw_seq_open);

static int raw_v4_seq_open(struct inode *inode, struct file *file)
{
- return raw_seq_open(file, &raw_v4_hashinfo, PF_INET);
+ return raw_seq_open(inode, file, &raw_v4_hashinfo, PF_INET);
}

static const struct file_operations raw_seq_fops = {
@@ -996,7 +1000,7 @@ static const struct file_operations raw_seq_fops = {
    .open = raw_v4_seq_open,
    .read = seq_read,

```

```
.llseek = seq_llseek,
- .release = seq_release_private,
+ .release = seq_release_net,
};

int __init raw_proc_init(void)
diff --git a/net/ipv6/raw.c b/net/ipv6/raw.c
index 6f20086..026fa91 100644
--- a/net/ipv6/raw.c
+++ b/net/ipv6/raw.c
@@ -1259,7 +1259,7 @@ static const struct seq_operations raw6_seq_ops = {

static int raw6_seq_open(struct inode *inode, struct file *file)
{
- return raw_seq_open(file, &raw_v6_hashinfo, PF_INET6);
+ return raw_seq_open(inode, file, &raw_v6_hashinfo, PF_INET6);
}

static const struct file_operations raw6_seq_fops = {
@@ -1267,7 +1267,7 @@ static const struct file_operations raw6_seq_fops = {
.open = raw6_seq_open,
.read = seq_read,
.llseek = seq_llseek,
- .release = seq_release_private,
+ .release = seq_release_net,
};

int __init raw6_proc_init(void)
--
1.5.3.4
```

Subject: Re: [PATCH net-2.6.25 2/4][NETNS][RAW]: Make /proc/net/raw(6) show per-namespace socket list.

Posted by [davem](#) on Mon, 14 Jan 2008 13:37:04 GMT

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

From: Pavel Emelyanov <xemul@openvz.org>

Date: Mon, 14 Jan 2008 16:08:42 +0300

> Pull the struct net pointer up to the showing functions

> to filter the sockets depending on their namespaces.

>

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

Applied.
