
Subject: [PATCH 0/3] [RAW]: proc output cleanups.
Posted by [den](#) on Thu, 31 Jan 2008 11:32:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

yesterday Adrian Bunk noticed, that the commit

commit 42a73808ed4f30b739eb52bcbb33a02fe62ceef5

Author: Pavel Emelyanov <xemul@openvz.org>

Date: Mon Nov 19 22:38:33 2007 -0800

is somewhat strange. Basically, the commit is obviously wrong as the content of the /proc/net/raw6 is incorrect due to it.

This series of patches fixes original problem and slightly cleanups the code around.

Signed-off-by: Denis V. Lunev <[den@openvz.org](#)>

Subject: [PATCH 1/3] [RAW]: Family check in the /proc/net/raw[6] is extra.
Posted by [den](#) on Thu, 31 Jan 2008 11:34:52 GMT

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

Different hashtables are used for IPv6 and IPv4 raw sockets, so no need to check the socket family in the iterator over hashtables. Clean this out.

Signed-off-by: Denis V. Lunev <[den@openvz.org](#)>

```
include/net/raw.h |  4 +---  
net/ipv4/raw.c   | 12 ++++++----  
net/ipv6/raw.c   |  2 +-  
3 files changed, 6 insertions(+), 12 deletions(-)
```

```
diff --git a/include/net/raw.h b/include/net/raw.h  
index cca81d8..c7ea7a2 100644  
--- a/include/net/raw.h  
+++ b/include/net/raw.h  
@@ -41,7 +41,6 @@ extern void raw_proc_exit(void);  
struct raw_iter_state {  
    struct seq_net_private p;  
    int bucket;  
-    unsigned short family;  
    struct raw_hashinfo *h;  
};  
  
@@ -49,8 +48,7 @@ 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 inode *ino, struct file *file, struct raw_hashinfo *h,
- unsigned short family);
+int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h);

#endif

diff --git a/net/ipv4/raw.c b/net/ipv4/raw.c
index f863c3d..507cbfe 100644
--- a/net/ipv4/raw.c
+++ b/net/ipv4/raw.c
@@ -862,8 +862,7 @@ 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_net == state->p.net &&
- sk->sk_family == state->family)
+ if (sk->sk_net == state->p.net)
     goto found;
 }
 sk = NULL;
@@ -879,8 +878,7 @@ static struct sock *raw_get_next(struct seq_file *seq, struct sock *sk)
 sk = sk_next(sk);
try_again:
 ;
- } while (sk && sk->sk_net != state->p.net &&
- sk->sk_family != state->family);
+ } while (sk && sk->sk_net != state->p.net);

 if (!sk && ++state->bucket < RAW_HTABLE_SIZE) {
     sk = sk_head(&state->h->ht[state->bucket]);
@@ -974,8 +972,7 @@ static const struct seq_operations raw_seq_ops = {
 .show = raw_seq_show,
};

-int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h,
- unsigned short family)
+int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h)
{
    int err;
    struct raw_iter_state *i;
@@ -987,14 +984,13 @@ int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo
*h,
    i = raw_seq_private((struct seq_file *)file->private_data);
    i->h = h;
- i->family = family;
    return 0;
}

```

```

}

EXPORT_SYMBOL_GPL(raw_seq_open);

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

static const struct file_operations raw_seq_fops = {
diff --git a/net/ipv6/raw.c b/net/ipv6/raw.c
index d61c63d..a2cf499 100644
--- a/net/ipv6/raw.c
+++ b/net/ipv6/raw.c
@@ -1262,7 +1262,7 @@ static const struct seq_operations raw6_seq_ops = {

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

static const struct file_operations raw6_seq_fops = {
-- 
1.5.3.rc5

```

Subject: [PATCH 2/3] [RAW]: Cleanup IPv4 raw_seq_show.

Posted by [den](#) on Thu, 31 Jan 2008 11:34:53 GMT

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

There is no need to use 128 bytes on the stack at all. Clean the code in the IPv6 style.

Signed-off-by: Denis V. Lunev <den@openvz.org>

```
---
net/ipv4/raw.c | 24 ++++++-----
1 files changed, 7 insertions(+), 17 deletions(-)
```

```
diff --git a/net/ipv4/raw.c b/net/ipv4/raw.c
index 507cbfe..830f19e 100644
--- a/net/ipv4/raw.c
+++ b/net/ipv4/raw.c
@@ -927,7 +927,7 @@ void raw_seq_stop(struct seq_file *seq, void *v)
}
EXPORT_SYMBOL_GPL(raw_seq_stop);

-static __inline__ char *get_raw_sock(struct sock *sp, char *tmpbuf, int i)
```

```

+static void raw_sock_seq_show(struct seq_file *seq, struct sock *sp, int i)
{
    struct inet_sock *inet = inet_sk(sp);
    __be32 dest = inet->daddr,
@@ -935,33 +935,23 @@ static __inline__ char *get_raw_sock(struct sock *sp, char *tmpbuf, int
i)
    __u16 destp = 0,
        srcp = inet->num;

- sprintf(tmpbuf, "%4d: %08X:%04X %08X:%04X"
+ seq_printf(seq, "%4d: %08X:%04X %08X:%04X"
" %02X %08X:%08X %02X:%08IX %08X %5d %8d %lu %d %p %d",
i, src, srcp, dest, destp, sp->sk_state,
atomic_read(&sp->sk_wmem_alloc),
atomic_read(&sp->sk_rmem_alloc),
0, 0L, 0, sock_i_uid(sp), 0, sock_i_ino(sp),
atomic_read(&sp->sk_refcnt), sp, atomic_read(&sp->sk_drops));
- return tmpbuf;
}

#define TMPSZ 128

-
static int raw_seq_show(struct seq_file *seq, void *v)
{
- char tmpbuf[TMPSZ+1];
-
if (v == SEQ_START_TOKEN)
- seq_printf(seq, "%-*s\n", TMPSZ-1,
-     " sl local_address rem_address st tx_queue "
-     "rx_queue tr tm->when retrnsmt uid timeout "
-     "inode drops");
- else {
- struct raw_iter_state *state = raw_seq_private(seq);
-
- seq_printf(seq, "%-*s\n", TMPSZ-1,
-     get_raw_sock(v, tmpbuf, state->bucket));
- }
+ seq_printf(seq, " sl local_address rem_address st tx_queue "
+     "rx_queue tr tm->when retrnsmt uid timeout "
+     "inode drops\n");
+ else
+ raw_sock_seq_show(seq, v, raw_seq_private(seq)->bucket);
return 0;
}

--
1.5.3.rc5

```

Subject: [PATCH 3/3] [RAW]: Wrong content of the /proc/net/raw6.
Posted by [den](#) on Thu, 31 Jan 2008 11:34:54 GMT

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

The address of IPv6 raw sockets was shown in the wrong format, from IPv4 ones.

The problem has been introduced by the

commit 42a73808ed4f30b739eb52bcbb33a02fe62ceef5

Author: Pavel Emelyanov <xemul@openvz.org>

Date: Mon Nov 19 22:38:33 2007 -0800

Thanks to Adrian Bunk who originally noticed the problem.

Signed-off-by: Denis V. Lunev <[den@openvz.org](#)>

```
include/net/raw.h |  3 ++
net/ipv4/raw.c   |  8 +++++-
net/ipv6/raw.c   |  2 ++
3 files changed, 7 insertions(+), 6 deletions(-)
```

```
diff --git a/include/net/raw.h b/include/net/raw.h
index c7ea7a2..1828f81 100644
--- a/include/net/raw.h
+++ b/include/net/raw.h
@@ -48,7 +48,8 @@ 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 inode *ino, struct file *file, struct raw_hashinfo *h);
+int raw_seq_open(struct inode *ino, struct file *file,
+ struct raw_hashinfo *h, const struct seq_operations *ops);

#endif
```

```
diff --git a/net/ipv4/raw.c b/net/ipv4/raw.c
index 830f19e..a3002fe 100644
--- a/net/ipv4/raw.c
+++ b/net/ipv4/raw.c
@@ -962,13 +962,13 @@ static const struct seq_operations raw_seq_ops = {
 .show = raw_seq_show,
 };

-int raw_seq_open(struct inode *ino, struct file *file, struct raw_hashinfo *h)
+int raw_seq_open(struct inode *ino, struct file *file,
+ struct raw_hashinfo *h, const struct seq_operations *ops)
{
 int err;
 struct raw_iter_state *i;

- err = seq_open_net(ino, file, &raw_seq_ops,
```

```

- sizeof(struct raw_iter_state));
+ err = seq_open_net(ino, file, ops, sizeof(struct raw_iter_state));
if (err < 0)
    return err;

@@ -980,7 +980,7 @@ EXPORT_SYMBOL_GPL(raw_seq_open);

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

static const struct file_operations raw_seq_fops = {
diff --git a/net/ipv6/raw.c b/net/ipv6/raw.c
index a2cf499..8897ccf 100644
--- a/net/ipv6/raw.c
+++ b/net/ipv6/raw.c
@@ -1262,7 +1262,7 @@ static const struct seq_operations raw6_seq_ops = {

static int raw6_seq_open(struct inode *inode, struct file *file)
{
- return raw_seq_open(inode, file, &raw_v6_hashinfo);
+ return raw_seq_open(inode, file, &raw_v6_hashinfo, &raw6_seq_ops);
}

static const struct file_operations raw6_seq_fops = {
--
```

1.5.3.rc5

Subject: Re: [PATCH 0/3] [RAW]: proc output cleanups.

Posted by [davem](#) on Thu, 31 Jan 2008 11:49:10 GMT

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

From: "Denis V. Lunev" <den@sw.ru>
Date: Thu, 31 Jan 2008 14:32:52 +0300

> yesterday Adrian Bunk noticed, that the commit
>
> commit 42a73808ed4f30b739eb52bcbb33a02fe62ceef5
> Author: Pavel Emelyanov <xemul@openvz.org>
> Date: Mon Nov 19 22:38:33 2007 -0800
>
> is somewhat strange. Basically, the commit is obviously wrong as the
> content of the /proc/net/raw6 is incorrect due to it.
>
> This series of patches fixes original problem and slightly cleanups the

> code around.
>
> Signed-off-by: Denis V. Lunev <den@openvz.org>

All applied, thanks a lot!
