
Subject: [PATCH 2/6][NET-2.6.24] Make core networking code use seq_open_private

Posted by [Pavel Emelianov](#) on Tue, 09 Oct 2007 15:55:28 GMT

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

This concerns the ipv4 and ipv6 code mostly, but also the netlink and unix sockets.

The netlink code is an example of how to use the __seq_open_private() call - it saves the net namespace on this private.

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

diff --git a/net/ipv4/arp.c b/net/ipv4/arp.c

index d824819..36d6798 100644

--- a/net/ipv4/arp.c

+++ b/net/ipv4/arp.c

@@ -1378,24 +1378,8 @@ static const struct seq_operations arp_s

```
static int arp_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct neigh_seq_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &arp_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &arp_seq_ops,
+ sizeof(struct neigh_seq_state));
}
```

static const struct file_operations arp_seq_fops = {

diff --git a/net/ipv4/fib_hash.c b/net/ipv4/fib_hash.c

index 9fafbee..527a6e0 100644

```

--- a/net/ipv4/fib_hash.c
+++ b/net/ipv4/fib_hash.c
@@ -1039,24 +1039,8 @@ static const struct seq_operations fib_s

```

```

static int fib_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct fib_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &fib_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &fib_seq_ops,
+ sizeof(struct fib_iter_state));
}

```

```

static const struct file_operations fib_seq_fops = {

```

```

diff --git a/net/ipv4/fib_trie.c b/net/ipv4/fib_trie.c

```

```

index be34bd5..81a8285 100644

```

```

--- a/net/ipv4/fib_trie.c
+++ b/net/ipv4/fib_trie.c
@@ -2379,25 +2379,8 @@ static const struct seq_operations fib_t

```

```

static int fib_trie_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct fib_trie_iter *s = kmalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &fib_trie_seq_ops);
- if (rc)
- goto out_kfree;
-

```

```

- seq    = file->private_data;
- seq->private = s;
- memset(s, 0, sizeof(*s));
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &fib_trie_seq_ops,
+ sizeof(struct fib_trie_iter));
}

```

```

static const struct file_operations fib_trie_fops = {
@@ -2500,25 +2483,8 @@ static const struct seq_operations fib_r

```

```

static int fib_route_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct fib_trie_iter *s = kmalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &fib_route_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq    = file->private_data;
- seq->private = s;
- memset(s, 0, sizeof(*s));
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &fib_route_seq_ops,
+ sizeof(struct fib_trie_iter));
}

```

```

static const struct file_operations fib_route_fops = {
diff --git a/net/ipv4/igmp.c b/net/ipv4/igmp.c
index 2b6e59c..7dbc282 100644
--- a/net/ipv4/igmp.c
+++ b/net/ipv4/igmp.c
@@ -2410,23 +2410,8 @@ static const struct seq_operations igmp_

```

```

static int igmp_mc_seq_open(struct inode *inode, struct file *file)

```

```

{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct igmp_mc_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
- rc = seq_open(file, &igmp_mc_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &igmp_mc_seq_ops,
+ sizeof(struct igmp_mc_iter_state));
}

static const struct file_operations igmp_mc_seq_fops = {
@@ -2584,23 +2569,8 @@ static const struct seq_operations igmp_

static int igmp_mcf_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct igmp_mcf_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
- rc = seq_open(file, &igmp_mcf_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &igmp_mcf_seq_ops,
+ sizeof(struct igmp_mcf_iter_state));
}

```

```

static const struct file_operations igmp_mcf_seq_fops = {
diff --git a/net/ipv4/ipmr.c b/net/ipv4/ipmr.c
index b8b4b49..37bb497 100644
--- a/net/ipv4/ipmr.c
+++ b/net/ipv4/ipmr.c
@@ -1714,26 +1714,8 @@ static const struct seq_operations ipmr_

```

```

static int ipmr_vif_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct ipmr_vif_iter *s = kmalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &ipmr_vif_seq_ops);
- if (rc)
- goto out_kfree;
-
- s->ct = 0;
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
-
+ return seq_open_private(file, &ipmr_vif_seq_ops,
+ sizeof(struct ipmr_vif_iter));
}

```

```

static const struct file_operations ipmr_vif_fops = {
@@ -1877,25 +1859,8 @@ static const struct seq_operations ipmr_

```

```

static int ipmr_mfc_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct ipmr_mfc_iter *s = kmalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &ipmr_mfc_seq_ops);
- if (rc)
- goto out_kfree;

```

```

-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
-
+ return seq_open_private(file, &ipmr_mfc_seq_ops,
+   sizeof(struct ipmr_mfc_iter));
}

static const struct file_operations ipmr_mfc_fops = {
diff --git a/net/ipv4/ipvs/ip_vs_ctl.c b/net/ipv4/ipvs/ip_vs_ctl.c
index 61d023d..7345fc2 100644
--- a/net/ipv4/ipvs/ip_vs_ctl.c
+++ b/net/ipv4/ipvs/ip_vs_ctl.c
@@ -1792,24 +1792,8 @@ static const struct seq_operations ip_vs

static int ip_vs_info_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct ip_vs_iter *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
-   goto out;
-
- rc = seq_open(file, &ip_vs_info_seq_ops);
- if (rc)
-   goto out_kfree;
-
- seq    = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &ip_vs_info_seq_ops,
+   sizeof(struct ip_vs_iter));
}

static const struct file_operations ip_vs_info_fops = {
diff --git a/net/ipv4/raw.c b/net/ipv4/raw.c
index 07070c7..3916fac 100644
--- a/net/ipv4/raw.c

```

```

+++ b/net/ipv4/raw.c
@@ -902,24 +902,8 @@ static const struct seq_operations raw_s

static int raw_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct raw_iter_state *s;
-
- s = kzalloc(sizeof(*s), GFP_KERNEL);
- if (!s)
- goto out;
- rc = seq_open(file, &raw_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &raw_seq_ops,
+ sizeof(struct raw_iter_state));
}

static const struct file_operations raw_seq_fops = {
diff --git a/net/ipv4/route.c b/net/ipv4/route.c
index 307e1f1..21b12de 100644
--- a/net/ipv4/route.c
+++ b/net/ipv4/route.c
@@ -375,23 +375,8 @@ static const struct seq_operations rt_ca

static int rt_cache_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct rt_cache_iter_state *s;
-
- s = kzalloc(sizeof(*s), GFP_KERNEL);
- if (!s)
- goto out;
- rc = seq_open(file, &rt_cache_seq_ops);
- if (rc)
- goto out_kfree;
- seq = file->private_data;
- seq->private = s;

```

```

-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &rt_cache_seq_ops,
+ sizeof(struct rt_cache_iter_state));
}

static const struct file_operations rt_cache_seq_fops = {
diff --git a/net/ipv6/addrconf.c b/net/ipv6/addrconf.c
index 6d5c3c2..8b2d760 100644
--- a/net/ipv6/addrconf.c
+++ b/net/ipv6/addrconf.c
@@ -2806,24 +2806,8 @@ static const struct seq_operations if6_s

static int if6_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct if6_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &if6_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &if6_seq_ops,
+ sizeof(struct if6_iter_state));
}

static const struct file_operations if6_fops = {
diff --git a/net/ipv6/anycast.c b/net/ipv6/anycast.c
index 5810852..f915c4d 100644
--- a/net/ipv6/anycast.c
+++ b/net/ipv6/anycast.c
@@ -549,24 +549,8 @@ static const struct seq_operations ac6_s

static int ac6_seq_open(struct inode *inode, struct file *file)

```



```

{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct ac6_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &ac6_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &ac6_seq_ops,
+ sizeof(struct ac6_iter_state));
}

static const struct file_operations ac6_seq_fops = {
diff --git a/net/ipv6/ip6_flowlabel.c b/net/ipv6/ip6_flowlabel.c
index 1791399..217d60f 100644
--- a/net/ipv6/ip6_flowlabel.c
+++ b/net/ipv6/ip6_flowlabel.c
@@ -658,24 +658,8 @@ static const struct seq_operations ip6fl

static int ip6fl_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct ip6fl_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
- goto out;
-
- rc = seq_open(file, &ip6fl_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:

```

```

- kfree(s);
- goto out;
+ return seq_open_private(file, &ip6fl_seq_ops,
+   sizeof(struct ip6fl_iter_state));
}

static const struct file_operations ip6fl_seq_fops = {
diff --git a/net/ipv6/mcast.c b/net/ipv6/mcast.c
index cc8d4e2..331d728 100644
--- a/net/ipv6/mcast.c
+++ b/net/ipv6/mcast.c
@@ -2426,24 +2426,8 @@ static const struct seq_operations igmp6

static int igmp6_mc_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct igmp6_mc_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
-   goto out;
-
- rc = seq_open(file, &igmp6_mc_seq_ops);
- if (rc)
-   goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &igmp6_mc_seq_ops,
+   sizeof(struct igmp6_mc_iter_state));
}

static const struct file_operations igmp6_mc_seq_fops = {
@@ -2600,24 +2584,8 @@ static const struct seq_operations igmp6

static int igmp6_mcf_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct igmp6_mcf_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
-
- if (!s)
-   goto out;

```

```

-
- rc = seq_open(file, &igmp6_mcf_seq_ops);
- if (rc)
- goto out_kfree;
-
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &igmp6_mcf_seq_ops,
+ sizeof(struct igmp6_mcf_iter_state));
}

static const struct file_operations igmp6_mcf_seq_fops = {
diff --git a/net/ipv6/raw.c b/net/ipv6/raw.c
index bdd0974..ca24ef1 100644
--- a/net/ipv6/raw.c
+++ b/net/ipv6/raw.c
@@ -1289,21 +1289,8 @@ static const struct seq_operations raw6_

```

```

static int raw6_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- struct raw6_iter_state *s = kzalloc(sizeof(*s), GFP_KERNEL);
- if (!s)
- goto out;
- rc = seq_open(file, &raw6_seq_ops);
- if (rc)
- goto out_kfree;
- seq = file->private_data;
- seq->private = s;
-out:
- return rc;
-out_kfree:
- kfree(s);
- goto out;
+ return seq_open_private(file, &raw6_seq_ops,
+ sizeof(struct raw6_iter_state));
}

```

```

static const struct file_operations raw6_seq_fops = {
diff --git a/net/netlink/af_netlink.c b/net/netlink/af_netlink.c
index 3ef3282..f934f54 100644
--- a/net/netlink/af_netlink.c

```

```

+++ b/net/netlink/af_netlink.c
@@ -1845,27 +1845,18 @@ static const struct seq_operations netli

static int netlink_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- struct nl_seq_iter *iter;
- int err;

- iter = kzalloc(sizeof(*iter), GFP_KERNEL);
+ iter = __seq_open_private(file, &netlink_seq_ops, sizeof(*iter));
  if (!iter)
    return -ENOMEM;

- err = seq_open(file, &netlink_seq_ops);
- if (err) {
-   kfree(iter);
-   return err;
- }
-
- seq = file->private_data;
- seq->private = iter;
  iter->net = get_proc_net(inode);
  if (!iter->net) {
    seq_release_private(inode, file);
    return -ENXIO;
  }
+
+ return 0;
}

```

```

diff --git a/net/unix/af_unix.c b/net/unix/af_unix.c
index 10e7312..2b57eaf 100644
--- a/net/unix/af_unix.c
+++ b/net/unix/af_unix.c
@@ -2086,25 +2086,7 @@ static const struct seq_operations unix_

static int unix_seq_open(struct inode *inode, struct file *file)
{
- struct seq_file *seq;
- int rc = -ENOMEM;
- int *iter = kmalloc(sizeof(int), GFP_KERNEL);
-
- if (!iter)
-   goto out;
-
- rc = seq_open(file, &unix_seq_ops);
- if (rc)

```

```
- goto out_kfree;
-
- seq    = file->private_data;
- seq->private = iter;
- *iter = 0;
-out:
- return rc;
-out_kfree:
- kfree(iter);
- goto out;
+ return seq_open_private(file, &unix_seq_ops, sizeof(int));
}
```

```
static const struct file_operations unix_seq_fops = {
```
