Subject: [PATCH -net 1/2] Convert /proc/net/ipv6_route to seq_file interface
Posted by Alexey Dobriyan on Tue, 06 Nov 2007 12:21:50 GMT
View Forum Message <> Reply to Message

This removes last proc_net_create() user. Kudos to Benjamin Thery and
Stephen Hemminger for comments on previous version.

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

```
---

 net/ipv6/route.c |   91 +++++++++++++++++-------------------------------------
 1 file changed, 29 insertions(+), 62 deletions(-)

--- a/net/ipv6/route.c
+++ b/net/ipv6/route.c
@@ -38,12 +38,8 @@
 #include <linux/in6.h>
 #include <linux/init.h>
 #include <linux/if_arp.h>
-
-#ifdef  CONFIG_PROC_FS
 #include <linux/proc_fs.h>
 #include <linux/seq_file.h>
-#endif
-
 #include <net/net_namespace.h>
 #include <net/snmp.h>
 #include <net/ipv6.h>
@@ -2288,71 +2284,50 @@ struct rt6_proc_arg

 static int rt6_info_route(struct rt6_info *rt, void *p_arg)
 {
- struct rt6_proc_arg *arg = (struct rt6_proc_arg *) p_arg;
+ struct seq_file *m = p_arg;

- if (arg->skip < arg->offset / RT6_INFO_LEN) {
-  arg->skip++;
-  return 0;
- }
-
- if (arg->len >= arg->length)
-  return 0;
-
- arg->len += sprintf(arg->buffer + arg->len,
-      NIP6_SEQFMT " %02x ",
-      NIP6(rt->rt6i_dst.addr),
-      rt->rt6i_dst.plen);
+ seq_printf(m, NIP6_SEQFMT " %02x ", NIP6(rt->rt6i_dst.addr),
```

```
+     rt->rt6i_dst.plen);

 #ifdef CONFIG_IPV6_SUBTREES
- arg->len += sprintf(arg->buffer + arg->len,
-       NIP6_SEQFMT " %02x ",
-       NIP6(rt->rt6i_src.addr),
-       rt->rt6i_src.plen);
+ seq_printf(m, NIP6_SEQFMT " %02x ", NIP6(rt->rt6i_src.addr),
+     rt->rt6i_src.plen);
 #else
- arg->len += sprintf(arg->buffer + arg->len,
-       "00000000000000000000000000000000 00 ");
+ seq_puts(m, "00000000000000000000000000000000 00 ");
 #endif

  if (rt->rt6i_nexthop) {
-  arg->len += sprintf(arg->buffer + arg->len,
-        NIP6_SEQFMT,
-        NIP6(*((struct in6_addr *)rt->rt6i_nexthop->primary_key)));
+  seq_printf(m, NIP6_SEQFMT,
+      NIP6(*((struct in6_addr *)rt->rt6i_nexthop->primary_key)));
  } else {
-  arg->len += sprintf(arg->buffer + arg->len,
-        "00000000000000000000000000000000");
+  seq_puts(m, "00000000000000000000000000000000");
  }
- arg->len += sprintf(arg->buffer + arg->len,
-       " %08x %08x %08x %08x %8s\n",
-       rt->rt6i_metric, atomic_read(&rt->u.dst.__refcnt),
-       rt->u.dst.__use, rt->rt6i_flags,
-       rt->rt6i_dev ? rt->rt6i_dev->name : "");
+ seq_printf(m, " %08x %08x %08x %08x %8s\n",
+     rt->rt6i_metric, atomic_read(&rt->u.dst.__refcnt),
+     rt->u.dst.__use, rt->rt6i_flags,
+     rt->rt6i_dev ? rt->rt6i_dev->name : "");
  return 0;
 }

-static int rt6_proc_info(char *buffer, char **start, off_t offset, int length)
+static int ipv6_route_show(struct seq_file *m, void *v)
 {
- struct rt6_proc_arg arg = {
-  .buffer = buffer,
-  .offset = offset,
-  .length = length,
- };
-
- fib6_clean_all(rt6_info_route, 0, &arg);
```

```
-
- *start = buffer;
- if (offset)
-  *start += offset % RT6_INFO_LEN;
-
- arg.len -= offset % RT6_INFO_LEN;
-
- if (arg.len > length)
-  arg.len = length;
- if (arg.len < 0)
-  arg.len = 0;
+ fib6_clean_all(rt6_info_route, 0, m);
+ return 0;
+}

- return arg.len;
+static int ipv6_route_open(struct inode *inode, struct file *file)
+{
+ return single_open(file, ipv6_route_show, NULL);
 }

+static const struct file_operations ipv6_route_proc_fops = {
+ .owner  = THIS_MODULE,
+ .open  = ipv6_route_open,
+ .read  = seq_read,
+ .llseek  = seq_lseek,
+ .release = single_release,
+};
+
 static int rt6_stats_seq_show(struct seq_file *seq, void *v)
 {
  seq_printf(seq, "%04x %04x %04x %04x %04x %04x %04x\n",
@@ -2489,22 +2464,14 @@ ctl_table ipv6_route_table[] = {

 void __init ip6_route_init(void)
 {
-#ifdef  CONFIG_PROC_FS
- struct proc_dir_entry *p;
-#endif
  ip6_dst_ops.kmem_cachep =
   kmem_cache_create("ip6_dst_cache", sizeof(struct rt6_info), 0,
       SLAB_HWCACHE_ALIGN|SLAB_PANIC, NULL);
  ip6_dst_blackhole_ops.kmem_cachep = ip6_dst_ops.kmem_cachep;

  fib6_init();
-#ifdef  CONFIG_PROC_FS
- p = proc_net_create(&init_net, "ipv6_route", 0, rt6_proc_info);
- if (p)
```

```
-  p->owner = THIS_MODULE;
-
+ proc_net_fops_create(&init_net, "ipv6_route", 0, &ipv6_route_proc_fops);
  proc_net_fops_create(&init_net, "rt6_stats", S_IRUGO, &rt6_stats_seq_fops);
-#endif
 #ifdef CONFIG_XFRM
  xfrm6_init();
 #endif
```
