
Subject: [PATCH net-next 3/9] ipv4: remove static flush_delay variable

Posted by [den](#) on Fri, 04 Jul 2008 13:17:06 GMT

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

flush delay is used as an external storage for net.ipv4.route.flush sysctl entry. It is write-only.

The ctl_table->data for this entry is used once. Fix this case to point to the stack to remove global variable. Do this to avoid additional variable on struct net in the next patch.

Possible race (as it was before) accessing this local variable is removed using flush_mutex.

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

net/ipv4/route.c | 11 +++++++--
1 files changed, 8 insertions(+), 3 deletions(-)

```
diff --git a/net/ipv4/route.c b/net/ipv4/route.c  
index cedc366..790de32 100644  
--- a/net/ipv4/route.c  
+++ b/net/ipv4/route.c  
@@ -2829,14 +2829,20 @@ void ip_rt_multicast_event(struct in_device *in_dev)  
}  
  
#ifdef CONFIG_SYSCTL  
-static int flush_delay;  
-  
 static int ipv4_sysctl_rtcache_flush(ctl_table *ctl, int write,  
     struct file *filp, void __user *buffer,  
     size_t *lenp, loff_t *ppos)  
{  
    if (write) {  
+    int flush_delay;  
+    static DEFINE_MUTEX(flush_mutex);  
+  
+    mutex_lock(&flush_mutex);  
+    ctl->data = &flush_delay;  
+    proc_dointvec(ctl, write, filp, buffer, lenp, ppos);  
+    ctl->data = NULL;  
+    mutex_unlock(&flush_mutex);  
+  
+    rt_cache_flush(&init_net, flush_delay);  
    return 0;  
 }  
@@ -2865,7 +2871,6 @@ ctl_table ipv4_route_table[] = {  
 {
```

```
.ctl_name = NET_IPV4_ROUTE_FLUSH,  
.procname = "flush",  
- .data = &flush_delay,  
.maxlen = sizeof(int),  
.mode = 0200,  
.proc_handler = &ipv4_sysctl_rtcache_flush,
```

--

1.5.3.rc5

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
