Subject: [PATCH 1/4] Un-define the IPTUNNEL_XMIT() macro
Posted by Pavel Emelianov on Fri, 09 Nov 2007 13:09:58 GMT
View Forum Message <> Reply to Message

This one is used in all the smth-to-ip tunnels we have and
looks ... not very good. Make this a regular function in the
tunnel4.ko module.

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

```
diff --git a/include/net/ipip.h b/include/net/ipip.h
index 7cdc914..bc8f4a1 100644
--- a/include/net/ipip.h
+++ b/include/net/ipip.h
@@ -25,23 +25,7 @@ struct ip_tunnel
  struct ip_tunnel_parm parms;
 };

-#define IPTUNNEL_XMIT() do {      \
- int err;       \
- int pkt_len = skb->len;      \
-         \
- skb->ip_summed = CHECKSUM_NONE;     \
- iph->tot_len = htons(skb->len);     \
- ip_select_ident(iph, &rt->u.dst, NULL);    \
- ip_send_check(iph);      \
-         \
- err = NF_HOOK(PF_INET, NF_IP_LOCAL_OUT, skb, NULL, rt->u.dst.dev, dst_output);\
- if (net_xmit_eval(err) == 0) {     \
-  stats->tx_bytes += pkt_len;    \
-  stats->tx_packets++;     \
- } else {       \
-  stats->tx_errors++;     \
-  stats->tx_aborted_errors++;    \
- }        \
-} while (0)
+void iptunnel_xmit(struct sk_buff *skb, struct rtable *rt,
+  struct net_device_stats *stats);

 #endif
diff --git a/net/ipv4/ip_gre.c b/net/ipv4/ip_gre.c
index 02b02a8..4fc19b3 100644
--- a/net/ipv4/ip_gre.c
+++ b/net/ipv4/ip_gre.c
@@ -882,7 +882,7 @@ static int ipgre_tunnel_xmit(struct sk_buff *skb, struct net_device *dev)
```

```
  nf_reset(skb);

- IPTUNNEL_XMIT();
+ iptunnel_xmit(skb, rt, stats);
  tunnel->recursion--;
  return 0;

diff --git a/net/ipv4/ipip.c b/net/ipv4/ipip.c
index 8c2b2b0..01f6d34 100644
--- a/net/ipv4/ipip.c
+++ b/net/ipv4/ipip.c
@@ -638,7 +638,7 @@ static int ipip_tunnel_xmit(struct sk_buff *skb, struct net_device *dev)

  nf_reset(skb);

- IPTUNNEL_XMIT();
+ iptunnel_xmit(skb, rt, stats);
  tunnel->recursion--;
  return 0;

diff --git a/net/ipv4/tunnel4.c b/net/ipv4/tunnel4.c
index a794a8c..c71617c 100644
--- a/net/ipv4/tunnel4.c
+++ b/net/ipv4/tunnel4.c
@@ -8,6 +8,7 @@
 #include <linux/mutex.h>
 #include <linux/netdevice.h>
 #include <linux/skbuff.h>
+#include <linux/netfilter_ipv4.h>
 #include <net/icmp.h>
 #include <net/ip.h>
 #include <net/protocol.h>
@@ -71,6 +72,32 @@ int xfrm4_tunnel_deregister(struct xfrm_tunnel *handler, unsigned short
family)

 EXPORT_SYMBOL(xfrm4_tunnel_deregister);

+void iptunnel_xmit(struct sk_buff *skb, struct rtable *rt,
+  struct net_device_stats *stats)
+{
+ int err;
+ struct iphdr *iph = ip_hdr(skb);
+ int pkt_len = skb->len;
+
+ skb->ip_summed = CHECKSUM_NONE;
+ iph->tot_len = htons(skb->len);
+ ip_select_ident(iph, &rt->u.dst, NULL);
+ ip_send_check(iph);
```

```
+
+ err = NF_HOOK(PF_INET, NF_IP_LOCAL_OUT, skb, NULL,
+   rt->u.dst.dev, dst_output);
+
+ if (net_xmit_eval(err) == 0) {
+  stats->tx_bytes += pkt_len;
+  stats->tx_packets++;
+ } else {
+  stats->tx_errors++;
+  stats->tx_aborted_errors++;
+ }
+}
+
+EXPORT_SYMBOL(iptunnel_xmit);
+
 static int tunnel4_rcv(struct sk_buff *skb)
 {
  struct xfrm_tunnel *handler;
diff --git a/net/ipv6/sit.c b/net/ipv6/sit.c
index 71433d2..49b30ca 100644
--- a/net/ipv6/sit.c
+++ b/net/ipv6/sit.c
@@ -579,7 +579,7 @@ static int ipip6_tunnel_xmit(struct sk_buff *skb, struct net_device *dev)

  nf_reset(skb);

- IPTUNNEL_XMIT();
+ iptunnel_xmit(skb, rt, stats);
  tunnel->recursion--;
  return 0;

--
1.5.3.4
```