
Subject: [PATCH RFC 27/31] net: Make the xfrm sysctls per network namespace.
Posted by [ebiederm](#) on Thu, 25 Jan 2007 19:00:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Eric W. Biederman <ebiederm@xmission.com> - unquoted

In particalure I moved:
/proc/sys/net/core/xfrm_aevent_etime
/proc/sys/net/core/xfrm_aevent_rseqth

Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

```
include/net/xfrm.h      |  4 ++--
net/core/sysctl_net_core.c | 37 ++++++-----
net/xfrm/xfrm_state.c   |  8 ++++---
net/xfrm/xfrm_user.c    | 10 ++++++---
4 files changed, 30 insertions(+), 29 deletions(-)
```

```
diff --git a/include/net/xfrm.h b/include/net/xfrm.h
index e476541..9b2e727 100644
--- a/include/net/xfrm.h
+++ b/include/net/xfrm.h
@@ -24,8 +24,8 @@
 MODULE_ALIAS("xfrm-mode-" __stringify(family) "-" __stringify(encap))
```

```
extern struct sock *xfrm_nl;
-extern u32 sysctl_xfrm_aevent_etime;
-extern u32 sysctl_xfrm_aevent_rseqth;
+DECLARE_PER_NET(u32, sysctl_xfrm_aevent_etime);
+DECLARE_PER_NET(u32, sysctl_xfrm_aevent_rseqth);
```

```
extern struct mutex xfrm_cfg_mutex;
```

```
diff --git a/net/core/sysctl_net_core.c b/net/core/sysctl_net_core.c
index 76f7a29..90f2a39 100644
--- a/net/core/sysctl_net_core.c
+++ b/net/core/sysctl_net_core.c
@@ -88,24 +88,6 @@ ctl_table core_table[] = {
     .mode = 0644,
     .proc_handler = &proc_dointvec
 },
-#ifdef CONFIG_XFRM
- {
-     .ctl_name = NET_CORE_AEVENT_ETIME,
-     .procname = "xfrm_aevent_etime",
-     .data = &sysctl_xfrm_aevent_etime,
-     .maxlen = sizeof(u32),
-     .mode = 0644,
```

```

- .proc_handler = &proc_dointvec
- },
- {
- .ctl_name = NET_CORE_AEVENT_RSEQTH,
- .procname = "xfrm_aevent_rseqth",
- .data = &sysctl_xfrm_aevent_rseqth,
- .maxlen = sizeof(u32),
- .mode = 0644,
- .proc_handler = &proc_dointvec
- },
-#endif /* CONFIG_XFRM */
-#endif /* CONFIG_NET */
{
    .ctl_name = NET_CORE_SOMAXCONN,
@@ -127,6 +109,23 @@ ctl_table core_table[] = {
};

DEFINE_PER_NET(struct ctl_table, multi_core_table[]) = {
- /* Stub for holding per network namespace sysctls */
-#ifdef CONFIG_XFRM
+ {
+ .ctl_name = NET_CORE_AEVENT_ETIME,
+ .procname = "xfrm_aevent_etime",
+ .data = &__per_net_base(sysctl_xfrm_aevent_etime),
+ .maxlen = sizeof(u32),
+ .mode = 0644,
+ .proc_handler = &proc_dointvec
+ },
+ {
+ .ctl_name = NET_CORE_AEVENT_RSEQTH,
+ .procname = "xfrm_aevent_rseqth",
+ .data = &__per_net_base(sysctl_xfrm_aevent_rseqth),
+ .maxlen = sizeof(u32),
+ .mode = 0644,
+ .proc_handler = &proc_dointvec
+ },
-#endif /* CONFIG_XFRM */
    {}
};
diff --git a/net/xfrm/xfrm_state.c b/net/xfrm/xfrm_state.c
index fdb08d9..3304a2d 100644
--- a/net/xfrm/xfrm_state.c
+++ b/net/xfrm/xfrm_state.c
@@ -27,11 +27,11 @@
struct sock *xfrm_nl;
EXPORT_SYMBOL(xfrm_nl);

-u32 sysctl_xfrm_aevent_etime = XFRM_AE_ETIME;

```

```

-EXPORT_SYMBOL(sysctl_xfrm_aevent_etime);
+DEFINE_PER_NET(u32, sysctl_xfrm_aevent_etime) = XFRM_AE_ETIME;
+EXPORT_PER_NET_SYMBOL(sysctl_xfrm_aevent_etime);

-u32 sysctl_xfrm_aevent_rseqth = XFRM_AE_SEQT_SIZE;
-EXPORT_SYMBOL(sysctl_xfrm_aevent_rseqth);
+DEFINE_PER_NET(u32, sysctl_xfrm_aevent_rseqth) = XFRM_AE_SEQT_SIZE;
+EXPORT_PER_NET_SYMBOL(sysctl_xfrm_aevent_rseqth);

/* Each xfrm_state may be linked to two tables:

diff --git a/net/xfrm/xfrm_user.c b/net/xfrm/xfrm_user.c
index 55affa7..15e962b 100644
--- a/net/xfrm/xfrm_user.c
+++ b/net/xfrm/xfrm_user.c
@@ -375,7 +375,8 @@ error:
     return err;
 }

-static struct xfrm_state *xfrm_state_construct(struct xfrm_usersa_info *p,
+static struct xfrm_state *xfrm_state_construct(net_t net,
+        struct xfrm_usersa_info *p,
+        struct rtattr **xfrma,
+        int *errp)
{
@@ -411,9 +412,9 @@ static struct xfrm_state *xfrm_state_construct(struct xfrm_usersa_info *p,
    goto error;

    x->km.seq = p->seq;
- x->replay_maxdiff = sysctl_xfrm_aevent_rseqth;
+ x->replay_maxdiff = per_net(sysctl_xfrm_aevent_rseqth, net);
    /* sysctl_xfrm_aevent_etime is in 100ms units */
- x->replay_maxage = (sysctl_xfrm_aevent_etime*HZ)/XFRM_AE_ETH_M;
+ x->replay_maxage = (per_net(sysctl_xfrm_aevent_etime, net)*HZ)/XFRM_AE_ETH_M;
    x->preplay.bitmap = 0;
    x->preplay.seq = x->replay.seq+x->replay_maxdiff;
    x->preplay.oseq = x->replay.oseq +x->replay_maxdiff;
@@ -437,6 +438,7 @@ error_no_put:
static int xfrm_add_sa(struct sk_buff *skb, struct nlmsg_hdr *nlh,
    struct rtattr **xfrma)
{
+ net_t net = skb->sk->sk_net;
    struct xfrm_usersa_info *p = NLMSG_DATA(nlh);
    struct xfrm_state *x;
    int err;
@@ -446,7 +448,7 @@ static int xfrm_add_sa(struct sk_buff *skb, struct nlmsg_hdr *nlh,
    if (err)
        return err;

```

```
- x = xfrm_state_construct(p, xfrma, &err);  
+ x = xfrm_state_construct(net, p, xfrma, &err);  
  if (!x)  
    return err;
```

--

1.4.4.1.g278f

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
