
Subject: [PATCH net-2.6.25 3/3]sysctl: make sysctl_somaxconn per-namespace

Posted by Pavel Emelianov on Fri, 07 Dec 2007 13:12:48 GMT

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

Just move the variable on the struct net and adjust its usage.

Others sysctls from sys.net.core table are more difficult to virtualize (i.e. make them per-namespace), but I'll look at them as well a bit later.

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

```
diff --git a/include/linux/socket.h b/include/linux/socket.h
index eb5bdd5..bd2b30a 100644
--- a/include/linux/socket.h
+++ b/include/linux/socket.h
@@ -24,7 +24,6 @@ struct __kernel_sockaddr_storage {
#include <linux/types.h> /* pid_t */
#include <linux/compiler.h> /* __user */
 
-extern int sysctl_somaxconn;
#endif CONFIG_PROC_FS
struct seq_file;
extern void socket_seq_show(struct seq_file *seq);
diff --git a/include/net/net_namespace.h b/include/net/net_namespace.h
index d593611..b62e31f 100644
--- a/include/net/net_namespace.h
+++ b/include/net/net_namespace.h
@@ -39,6 +39,7 @@ struct net {
 
 /* core sysctls */
 struct ctl_table_header *sysctl_core_hdr;
+int sysctl_somaxconn;
 
 /* List of all packet sockets. */
 rwlock_t packet_sklist_lock;
diff --git a/net/core/sysctl_net_core.c b/net/core/sysctl_net_core.c
index dc4cf7d..130338f 100644
--- a/net/core/sysctl_net_core.c
+++ b/net/core/sysctl_net_core.c
@@ -127,7 +127,7 @@ static struct ctl_table net_core_table[] = {
 {
 .ctl_name = NET_CORE_SOMAXCONN,
 .procname = "somaxconn",
- .data = &sysctl_somaxconn,
```

```

+ .data = &init_net.sysctl_somaxconn,
  . maxlen = sizeof(int),
  . mode = 0644,
  .proc_handler = &proc_dointvec
@@ -161,6 +161,8 @@ static __net_init int sysctl_core_net_init(struct net *net)
{
    struct ctl_table *tbl, *tmp;

+ net->sysctl_somaxconn = SOMAXCONN;
+
    tbl = net_core_table;
    if (net != &init_net) {
        tbl = kmemdup(tbl, sizeof(net_core_table), GFP_KERNEL);
diff --git a/net/socket.c b/net/socket.c
index 9ebca5c..7651de0 100644
--- a/net/socket.c
+++ b/net/socket.c
@@ -1365,17 +1365,17 @@ asmlinkage long sys_bind(int fd, struct sockaddr __user *umyaddr,
int addrlen)
    * ready for listening.
 */

```

```

-int sysctl_somaxconn __read_mostly = SOMAXCONN;
-
asmlinkage long sys_listen(int fd, int backlog)
{
    struct socket *sock;
    int err, fput_needed;
+ int somaxconn;

    sock = sockfd_lookup_light(fd, &err, &fput_needed);
    if (sock) {
- if ((unsigned)backlog > sysctl_somaxconn)
-     backlog = sysctl_somaxconn;
+ somaxconn = sock->sk->sk_net->sysctl_somaxconn;
+ if ((unsigned)backlog > somaxconn)
+     backlog = somaxconn;

    err = security_socket_listen(sock, backlog);
    if (!err)
--
```

1.5.3.4
