
Subject: [PATCH RFC 28/31] net: Make the SOMAXCONN sysctl per network namespace

Posted by [ebiederm](#) on Thu, 25 Jan 2007 19:00:30 GMT

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

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

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

```
include/linux/socket.h    |   3 ++-
net/core/sysctl_net_core.c |  16 ++++++-----
net/socket.c              |   7 +++++-
3 files changed, 14 insertions(+), 12 deletions(-)
```

diff --git a/include/linux/socket.h b/include/linux/socket.h

index 92cd38e..aa159ea 100644

--- a/include/linux/socket.h

+++ b/include/linux/socket.h

```
@ @ -23,8 +23,9 @ @ struct __kernel_sockaddr_storage {
#include <linux/uio.h> /* iovec support */
#include <linux/types.h> /* pid_t */
#include <linux/compiler.h> /* __user */
+#include <linux/net_namespace_type.h>
```

```
-extern int sysctl_somaxconn;
```

```
+DECLARE_PER_NET(int, sysctl_somaxconn);
```

```
#ifdef CONFIG_PROC_FS
```

```
struct seq_file;
```

```
extern void socket_seq_show(struct seq_file *seq);
```

diff --git a/net/core/sysctl_net_core.c b/net/core/sysctl_net_core.c

index 90f2a39..14eca68 100644

--- a/net/core/sysctl_net_core.c

+++ b/net/core/sysctl_net_core.c

```
@ @ -90,14 +90,6 @ @ ctl_table core_table[] = {
},
```

```
#endif /* CONFIG_NET */
```

```
{
```

```
- .ctl_name = NET_CORE_SOMAXCONN,
```

```
- .procname = "somaxconn",
```

```
- .data = &sysctl_somaxconn,
```

```
- .maxlen = sizeof(int),
```

```
- .mode = 0644,
```

```
- .proc_handler = &proc_dointvec
```

```
- },
```

```
- {
```

```
    .ctl_name = NET_CORE_BUDGET,
```

```
    .procname = "netdev_budget",
```

```
    .data = &netdev_budget,
```

```

@@ -127,5 +119,13 @@ DEFINE_PER_NET(struct ctl_table, multi_core_table[]) = {
    .proc_handler = &proc_dointvec
    },
#endif /* CONFIG_XFRM */
+ {
+ .ctl_name = NET_CORE_SOMAXCONN,
+ .procname = "somaxconn",
+ .data = &__per_net_base(sysctl_somaxconn),
+ .maxlen = sizeof(int),
+ .mode = 0644,
+ .proc_handler = &proc_dointvec
+ },
    {}
};
diff --git a/net/socket.c b/net/socket.c
index 7371654..ab2aeea 100644
--- a/net/socket.c
+++ b/net/socket.c
@@ -1305,7 +1305,7 @@ asmlinkage long sys_bind(int fd, struct sockaddr __user *umyaddr, int
addrlen)
    * ready for listening.
    */

-int sysctl_somaxconn __read_mostly = SOMAXCONN;
+DEFINE_PER_NET(int, sysctl_somaxconn) = SOMAXCONN;

asmlinkage long sys_listen(int fd, int backlog)
{
@@ -1314,8 +1314,9 @@ asmlinkage long sys_listen(int fd, int backlog)

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

    err = security_socket_listen(sock, backlog);
    if (!err)
--
1.4.4.1.g278f

```

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