
Subject: [PATCH 2/4 net-2.6.25][UNIX] Move the sysctl_unix_max_dgram_qlen on struct net

Posted by Pavel Emelianov on Fri, 30 Nov 2007 16:29:03 GMT

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

This will make all the sub-namespaces always use the default value (10) and leave the tuning via sysctl to the init namespace only.

Per-namespace tuning is coming.

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

```
diff --git a/include/net/net_namespace.h b/include/net/net_namespace.h
index 235214c..b0cf075 100644
--- a/include/net/net_namespace.h
+++ b/include/net/net_namespace.h
@@ -38,6 +38,9 @@ struct net {
 /* List of all packet sockets. */
 rwlock_t packet_sklist_lock;
 struct hlist_head packet_sklist;
+
+/* unix sockets */
+int sysctl_unix_max_dgram_qlen;
};

#endif CONFIG_NET
diff --git a/net/unix/af_unix.c b/net/unix/af_unix.c
index a0aa6d3..73620d6 100644
--- a/net/unix/af_unix.c
+++ b/net/unix/af_unix.c
@@ -117,8 +117,6 @@ 
#include <net/checksum.h>
#include <linux/security.h>

-int sysctl_unix_max_dgram_qlen __read_mostly = 10;
-
 static struct hlist_head unix_socket_table[UNIX_HASH_SIZE + 1];
 static DEFINE_SPINLOCK(unix_table_lock);
 static atomic_t unix_nr_socks = ATOMIC_INIT(0);
@@ -594,7 +592,7 @@ static struct sock * unix_create1(struct net *net, struct socket *sock)
     &af_unix_sk_receive_queue_lock_key);

 sk->sk_write_space = unix_write_space;
- sk->sk_max_ack_backlog = sysctl_unix_max_dgram_qlen;
+ sk->sk_max_ack_backlog = net->sysctl_unix_max_dgram_qlen;
```

```
sk->sk_destruct = unix_sock_destructor;
u = unix_sk(sk);
u->dentry = NULL;
@@ -2140,6 +2138,8 @@ static int unix_net_init(struct net *net)
{
int error = -ENOMEM;

+ net->sysctl_unix_max_dgram_qlen = 10;
+
#ifndef CONFIG_PROC_FS
if (!proc_net_fops_create(net, "unix", 0, &unix_seq_fops))
    goto out;
diff --git a/net/unix/sysctl_net_unix.c b/net/unix/sysctl_net_unix.c
index b2e0407..c46ceco 100644
--- a/net/unix/sysctl_net_unix.c
+++ b/net/unix/sysctl_net_unix.c
@@ -18,7 +18,7 @@ static ctl_table unix_table[] = {
{
.ctl_name = NET_UNIX_MAX_DGRAM_QLEN,
.procname = "max_dgram_qlen",
-.data = &sysctl_unix_max_dgram_qlen,
+.data = &init_net.sysctl_unix_max_dgram_qlen,
.maxlen = sizeof(int),
.mode = 0644,
.proc_handler = &proc_dointvec
--
```

1.5.3.4

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
