

---

Subject: [PATCH 4/4] sysctl: Fix sys\_sysctl interface of ipc sysctls

Posted by [ebiederm](#) on Mon, 27 Nov 2006 05:05:11 GMT

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

---

Currently there is a regression and the ipc sysctls  
don't show up in the binary sysctl namespace.

This patch adds sysctl\_ipc\_data to read data/write from the  
appropriate namespace and deliver it in the expected manner.

Signed-off-by: Eric W. Biederman <[ebiederm@xmission.com](mailto:ebiederm@xmission.com)>

---

```
kernel/sysctl.c | 50 ++++++++++++++++++++++++++++++++++++++++
1 files changed, 50 insertions(+), 0 deletions(-)

diff --git a/kernel/sysctl.c b/kernel/sysctl.c
index 638aa14..24c2ca8 100644
--- a/kernel/sysctl.c
+++ b/kernel/sysctl.c
@@ -142,6 +142,10 @@ static int sysctl_uts_string(ctl_table *
    void __user *oldval, size_t __user *oldlenp,
    void __user *newval, size_t newlen, void **context);

+static int sysctl_ipc_data(ctl_table *table, int __user *name, int nlen,
+  void __user *oldval, size_t __user *oldlenp,
+  void __user *newval, size_t newlen, void **context);
+
#endif CONFIG_PROC_SYSCTL
static int proc_do_cad_pid(ctl_table *table, int write, struct file *filp,
    void __user *buffer, size_t *lenp, loff_t *ppos);
@@ -475,6 +479,7 @@ #ifdef CONFIG_SYSVIPC
    .maxlen = sizeof (init_ipc_ns.shm_ctlmax),
    .mode = 0644,
    .proc_handler = &proc_ipc_doulongvec_minmax,
+   .strategy = sysctl_ipc_data,
},
{
    .ctl_name = KERN_SHMALL,
@@ -483,6 +488,7 @@ #ifdef CONFIG_SYSVIPC
    .maxlen = sizeof (init_ipc_ns.shm_ctlall),
    .mode = 0644,
    .proc_handler = &proc_ipc_doulongvec_minmax,
+   .strategy = sysctl_ipc_data,
},
{
    .ctl_name = KERN_SHMMNI,
@@ -491,6 +497,7 @@ #ifdef CONFIG_SYSVIPC
    .maxlen = sizeof (init_ipc_ns.shm_ctlmni),
```

```

.mode = 0644,
.proc_handler = &proc_ipc_dointvec,
+ .strategy = sysctl_ipc_data,
},
{
.ctl_name = KERN_MSGMAX,
@@ -499,6 +506,7 @@ @ @ #ifdef CONFIG_SYSVIPC
 maxlen = sizeof (init_ipc_ns.msg_ctlmax),
.mode = 0644,
.proc_handler = &proc_ipc_dointvec,
+ .strategy = sysctl_ipc_data,
},
{
.ctl_name = KERN_MSGMNI,
@@ -507,6 +515,7 @@ @ @ #ifdef CONFIG_SYSVIPC
 maxlen = sizeof (init_ipc_ns.msg_ctlmni),
.mode = 0644,
.proc_handler = &proc_ipc_dointvec,
+ .strategy = sysctl_ipc_data,
},
{
.ctl_name = KERN_MSGMNB,
@@ -515,6 +524,7 @@ @ @ #ifdef CONFIG_SYSVIPC
 maxlen = sizeof (init_ipc_ns.msg_ctlmnb),
.mode = 0644,
.proc_handler = &proc_ipc_dointvec,
+ .strategy = sysctl_ipc_data,
},
{
.ctl_name = KERN_SEM,
@@ -523,6 +533,7 @@ @ @ #ifdef CONFIG_SYSVIPC
 maxlen = 4*sizeof (int),
.mode = 0644,
.proc_handler = &proc_ipc_dointvec,
+ .strategy = sysctl_ipc_data,
},
#endif
#endif CONFIG_MAGIC_SYSRQ
@@ -2613,6 +2624,45 @@ static int sysctl_uts_string(ctl_table *
 return r;
}

/* The generic sysctl ipc data routine. */
+static int sysctl_ipc_data(ctl_table *table, int __user *name, int nlen,
+ void __user *oldval, size_t __user *oldlenp,
+ void __user *newval, size_t newlen, void **context)
+{
+ size_t len;

```

```
+ void *data;
+
+ /* Get out of I don't have a variable */
+ if (!table->data || !table->maxlen)
+     return -ENOTDIR;
+
+ data = get_ipc(table, 1);
+ if (!data)
+     return -ENOTDIR;
+
+ if (oldval && oldlenp) {
+     if (get_user(len, oldlenp))
+         return -EFAULT;
+     if (len) {
+         if (len > table->maxlen)
+             len = table->maxlen;
+         if (copy_to_user(oldval, data, len))
+             return -EFAULT;
+         if (put_user(len, oldlenp))
+             return -EFAULT;
+     }
+ }
+
+ if (newval && newlen) {
+     if (newlen > table->maxlen)
+         newlen = table->maxlen;
+
+     if (copy_from_user(data, newval, newlen))
+         return -EFAULT;
+ }
+ return 1;
+}
+
#endif /* CONFIG_SYSCTL_SYSCALL */
```

--  
1.4.2.rc3.g7e18e-dirty

---

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

---