
Subject: [PATCH 6/6] IPC namespace - sysctls
Posted by [Kirill Korotaev](#) on Fri, 09 Jun 2006 15:11:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sysctl tweaks for IPC namespace

Signed-Off-By: Pavel Emelianiov <xemul@openvz.org>

Signed-Off-By: Kirill Korotaev <dev@openvz.org>

```
--- ./kernel/sysctl.c.ipcns 2006-06-06 14:47:59.000000000 +0400
+++ ./kernel/sysctl.c 2006-06-08 15:29:58.000000000 +0400
@@ -104,13 +104,8 @@ extern char modprobe_path[];
extern int sg_big_buff;
#endif
#ifdef CONFIG_SYSVIPC
-extern size_t shm_ctlmax;
-extern size_t shm_ctlall;
-extern int shm_ctlmni;
-extern int msg_ctlmax;
-extern int msg_ctlmnb;
-extern int msg_ctlmni;
-extern int sem_ctls[];
+static int proc_do_ipc_string(ctl_table *table, int write, struct file *filp,
+ void __user *buffer, size_t *lenp, loff_t *ppos);
#endif

#ifdef __sparc__
@@ -510,58 +505,58 @@ static ctl_table kern_table[] = {
{
    .ctl_name = KERN_SHMMAX,
    .procname = "shmmax",
-   .data = &shm_ctlmax,
+   .data = NULL,
    .maxlen = sizeof (size_t),
    .mode = 0644,
-   .proc_handler = &proc_doulongvec_minmax,
+   .proc_handler = &proc_do_ipc_string,
},
{
    .ctl_name = KERN_SHMALL,
    .procname = "shmall",
-   .data = &shm_ctlall,
+   .data = NULL,
    .maxlen = sizeof (size_t),
    .mode = 0644,
-   .proc_handler = &proc_doulongvec_minmax,
+   .proc_handler = &proc_do_ipc_string,
},
}
```

```

{
    .ctl_name = KERN_SHMMNI,
    .procname = "shmmni",
-   .data = &shm_ctlmni,
+   .data = NULL,
    .maxlen = sizeof (int),
    .mode = 0644,
-   .proc_handler = &proc_dointvec,
+   .proc_handler = &proc_do_ipc_string,
},
{
    .ctl_name = KERN_MSGMAX,
    .procname = "msgmax",
-   .data = &msg_ctlmax,
+   .data = NULL,
    .maxlen = sizeof (int),
    .mode = 0644,
-   .proc_handler = &proc_dointvec,
+   .proc_handler = &proc_do_ipc_string,
},
{
    .ctl_name = KERN_MSGMNI,
    .procname = "msgmni",
-   .data = &msg_ctlmni,
+   .data = NULL,
    .maxlen = sizeof (int),
    .mode = 0644,
-   .proc_handler = &proc_dointvec,
+   .proc_handler = &proc_do_ipc_string,
},
{
    .ctl_name = KERN_MSGMNB,
    .procname = "msgmnb",
-   .data = &msg_ctlmnb,
+   .data = NULL,
    .maxlen = sizeof (int),
    .mode = 0644,
-   .proc_handler = &proc_dointvec,
+   .proc_handler = &proc_do_ipc_string,
},
{
    .ctl_name = KERN_SEM,
    .procname = "sem",
-   .data = &sem_ctls,
+   .data = NULL,
    .maxlen = 4*sizeof (int),
    .mode = 0644,
-   .proc_handler = &proc_dointvec,

```

```

+ .proc_handler = &proc_do_ipc_string,
+ },
+ #endif
+ #ifdef CONFIG_MAGIC_SYSRQ
+ @@ -1866,8 +1861,9 @@ static int do_proc_dointvec_conv(int *ne
+     return 0;
+ }

-static int do_proc_dointvec(ctl_table *table, int write, struct file *filp,
- void __user *buffer, size_t *lenp, loff_t *ppos,
+static int __do_proc_dointvec(void *tbl_data, ctl_table *table,
+ int write, struct file *filp, void __user *buffer,
+ size_t *lenp, loff_t *ppos,
+     int (*conv)(int *negp, unsigned long *lvalp, int *valp,
+         int write, void *data),
+     void *data)
+@@ -1880,13 +1876,13 @@ static int do_proc_dointvec(ctl_table *t
+     char buf[TMPBUFLLEN], *p;
+     char __user *s = buffer;

- if (!table->data || !table->maxlen || !*lenp ||
+ if (!tbl_data || !table->maxlen || !*lenp ||
+     (*ppos && !write)) {
+     *lenp = 0;
+     return 0;
+ }

- i = (int *) table->data;
+ i = (int *) tbl_data;
+ vleft = table->maxlen / sizeof(*i);
+ left = *lenp;

@@ -1975,6 +1971,16 @@ static int do_proc_dointvec(ctl_table *t
+ #undef TMPBUFLLEN
+ }

+static int do_proc_dointvec(ctl_table *table, int write, struct file *filp,
+ void __user *buffer, size_t *lenp, loff_t *ppos,
+ int (*conv)(int *negp, unsigned long *lvalp, int *valp,
+     int write, void *data),
+ void *data)
+{
+ return __do_proc_dointvec(table->data, table, write, filp,
+     buffer, lenp, ppos, conv, data);
+}
+
+/**
+ * proc_dointvec - read a vector of integers

```

```

* @table: the sysctl table
@@ -2108,7 +2114,7 @@ int proc_dointvec_minmax(ctl_table *tbl
    do_proc_dointvec_minmax_conv, &param);
}

-static int do_proc_doulongvec_minmax(ctl_table *table, int write,
+static int __do_proc_doulongvec_minmax(void *data, ctl_table *table, int write,
    struct file *filp,
    void __user *buffer,
    size_t *lenp, loff_t *ppos,
@@ -2122,13 +2128,13 @@ static int do_proc_doulongvec_minmax(ctl
    char buf[TMPBUFLen], *p;
    char __user *s = buffer;

- if (!table->data || !table->maxlen || !*lenp ||
+ if (!data || !table->maxlen || !*lenp ||
    (*ppos && !write)) {
    *lenp = 0;
    return 0;
}

- i = (unsigned long *) table->data;
+ i = (unsigned long *) data;
    min = (unsigned long *) table->extra1;
    max = (unsigned long *) table->extra2;
    vleft = table->maxlen / sizeof(unsigned long);
@@ -2213,6 +2219,17 @@ static int do_proc_doulongvec_minmax(ctl
#undef TMPBUFLen
}

+static int do_proc_doulongvec_minmax(ctl_table *table, int write,
+    struct file *filp,
+    void __user *buffer,
+    size_t *lenp, loff_t *ppos,
+    unsigned long convmul,
+    unsigned long convdiv)
+{
+ return __do_proc_doulongvec_minmax(table->data, table, write,
+    filp, buffer, lenp, ppos, convmul, convdiv);
+}
+
+/**
+ * proc_doulongvec_minmax - read a vector of long integers with min/max values
+ * @table: the sysctl table
@@ -2401,6 +2418,49 @@ int proc_dointvec_ms_jiffies(ctl_table *
    do_proc_dointvec_ms_jiffies_conv, NULL);
}

```

```

+ #ifdef CONFIG_SYSVIP
+ static int proc_do_ipc_string(ctl_table *table, int write, struct file *filp,
+ void __user *buffer, size_t *lenp, loff_t *ppos)
+ {
+ void *data;
+ struct ipc_namespace *ns;
+
+ ns = current->nsproxy->ipc_ns;
+
+ switch (table->ctl_name) {
+ case KERN_SHMMAX:
+ data = &ns->shm_ctlmax;
+ goto proc_minmax;
+ case KERN_SHMALL:
+ data = &ns->shm_ctlall;
+ goto proc_minmax;
+ case KERN_SHMMNI:
+ data = &ns->shm_ctlmni;
+ break;
+ case KERN_MSGMAX:
+ data = &ns->msg_ctlmax;
+ break;
+ case KERN_MSGMNI:
+ data = &ns->msg_ctlmni;
+ break;
+ case KERN_MSGMNB:
+ data = &ns->msg_ctlmnb;
+ break;
+ case KERN_SEM:
+ data = &ns->sem_ctls;
+ break;
+ default:
+ return -EINVAL;
+ }
+
+ return __do_proc_dointvec(data, table, write, filp, buffer,
+ lenp, ppos, NULL, NULL);
+ proc_minmax:
+ return __do_proc_doulongvec_minmax(data, table, write, filp, buffer,
+ lenp, ppos, 1l, 1l);
+ }
+ #endif
+
+ #else /* CONFIG_PROC_FS */

int proc_dostring(ctl_table *table, int write, struct file *filp,

```
