

---

Subject: [RFC][PATCH v2 7/7] Do not recompute msgmni anymore if explicitly set by user

Posted by [Nadia Derby](#) on Thu, 31 Jan 2008 13:40:25 GMT

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

---

[PATCH 07/07]

This patch makes msgmni not recomputed anymore upon ipc namespace creation / removal or memory add/remove, as soon as it has been set from userland.

As soon as msgmni is explicitly set via procfs or sysctl(), the associated callback routine is unregistered from the ipc namespace notifier chain.

Signed-off-by: Nadia Derby <[Nadia.Derbey@bull.net](mailto:Nadia.Derbey@bull.net)>

---  
ipc/ipc\_sysctl.c | 43 +++  
1 file changed, 41 insertions(+), 2 deletions(-)

Index: linux-2.6.24/ipc/ipc\_sysctl.c

=====

--- linux-2.6.24.orig/ipc/ipc\_sysctl.c 2008-01-29 16:55:04.000000000 +0100

+++ linux-2.6.24/ipc/ipc\_sysctl.c 2008-01-31 13:13:14.000000000 +0100

@@ -34,6 +34,24 @@ static int proc\_ipc\_dointvec(ctl\_table \*  
return proc\_dointvec(&ipc\_table, write, filp, buffer, lenp, ppos);  
}

+static int proc\_ipc\_callback\_dointvec(ctl\_table \*table, int write,  
+ struct file \*filp, void \_\_user \*buffer, size\_t \*lenp, loff\_t \*ppos)

+{

+ size\_t lenp\_bef = \*lenp;

+ int rc;

+

+ rc = proc\_ipc\_dointvec(table, write, filp, buffer, lenp, ppos);

+

+ if (write && !rc && lenp\_bef == \*lenp)

+ /\*

+ \* Tunable has successfully been changed from userland:

+ \* disable its automatic recomputing.

+ \*/

+ unregister\_ipcns\_notifier(current->nsproxy->ipc\_ns);

+

+ return rc;

+}

+

static int proc\_ipc\_doulongvec\_minmax(ctl\_table \*table, int write,  
struct file \*filp, void \_\_user \*buffer, size\_t \*lenp, loff\_t \*ppos)

```

{
@@ -48,6 +66,7 @@ static int proc_ipc_doulongvec_minmax(ct
#else
#define proc_ipc_doulongvec_minmax NULL
#define proc_ipc_dointvec NULL
+#define proc_ipc_callback_dointvec NULL
#endif

#ifdef CONFIG_SYSCTL_SYSCALL
@@ -89,8 +108,28 @@ static int sysctl_ipc_data(ctl_table *ta
}
return 1;
}
+
+static int sysctl_ipc_registered_data(ctl_table *table, int __user *name,
+ int nlen, void __user *oldval, size_t __user *oldlenp,
+ void __user *newval, size_t newlen)
+{
+ int rc;
+
+ rc = sysctl_ipc_data(table, name, nlen, oldval, oldlenp, newval,
+ newlen);
+
+ if (newval && newlen && rc > 0)
+ /*
+ * Tunable has successfully been changed from userland:
+ * disable its automatic recomputing.
+ */
+ unregister_ipcns_notifier(current->nsproxy->ipc_ns);
+
+ return rc;
+}
#else
#define sysctl_ipc_data NULL
+#define sysctl_ipc_registered_data NULL
#endif

static struct ctl_table ipc_kern_table[] = {
@@ -136,8 +175,8 @@ static struct ctl_table ipc_kern_table[]
.data = &init_ipc_ns.msg_ctlmni,
.maxlen = sizeof (init_ipc_ns.msg_ctlmni),
.mode = 0644,
- .proc_handler = proc_ipc_dointvec,
- .strategy = sysctl_ipc_data,
+ .proc_handler = proc_ipc_callback_dointvec,
+ .strategy = sysctl_ipc_registered_data,
},
{

```

.ctl\_name = KERN\_MSGMNB,

--

---

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

---