
Subject: [RFC][PATCH 5/5] Subject: mqueue namespace: adapt sysctl
Posted by [Dave Hansen](#) on Thu, 10 Jul 2008 22:30:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Cedric Le Goater <clg@fr.ibm.com>

Largely inspired from ipc/ipc_sysctl.c. This patch isolates the mqueue sysctl stuff in its own file.

Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>

```
linux-2.6.git-dave/include/linux/mq_namespace.h | 10 ++
linux-2.6.git-dave/init/Kconfig | 6 +
linux-2.6.git-dave/ipc/Makefile | 1
linux-2.6.git-dave/ipc/mq_sysctl.c | 117 ++++++=====
linux-2.6.git-dave/ipc/mqueue.c | 56 -----
5 files changed, 135 insertions(+), 55 deletions(-)
```

```
diff -puN include/linux/mq_namespace.h~mq_namespace-fix-sysctl
include/linux/mq_namespace.h
--- linux-2.6.git/include/linux/mq_namespace.h~mq_namespace-fix-sysctl 2008-06-24
12:03:19.000000000 -0700
+++ linux-2.6.git-dave/include/linux/mq_namespace.h 2008-06-24 12:03:19.000000000 -0700
@@ -67,4 +67,14 @@ static inline void put_mq_ns(struct mq_n
#endif /* CONFIG_POSIX_MQUEUE */

+ifdef CONFIG_POSIX_MQUEUE_SYSCTL
+struct ctl_table_header;
+extern struct ctl_table_header *mq_register_sysctl_table(void);
+#else
+static inline struct ctl_table_header *mq_register_sysctl_table(void)
+{
+ return NULL;
+}
#endif /* CONFIG_POSIX_MQUEUE_SYSCTL */
+
#endif /* _LINUX_MQ_H */
diff -puN init/Kconfig~mq_namespace-fix-sysctl init/Kconfig
--- linux-2.6.git/init/Kconfig~mq_namespace-fix-sysctl 2008-06-24 12:03:19.000000000 -0700
+++ linux-2.6.git-dave/init/Kconfig 2008-06-24 12:03:19.000000000 -0700
@@ -148,6 +148,12 @@ config POSIX_MQUEUE
```

If unsure, say Y.

```
+config POSIX_MQUEUE_SYSCTL
+bool
```

```

+ depends on POSIX_MQUEUE
+ depends on SYSCTL
+ default y
+
config BSD_PROCESS_ACCT
    bool "BSD Process Accounting"
    help
diff -puN ipc/Makefile~mq_namespace-fix-sysctl ipc/Makefile
--- linux-2.6.git/ipc/Makefile~mq_namespace-fix-sysctl 2008-06-24 12:03:19.000000000 -0700
+++ linux-2.6.git-dave/ipc/Makefile 2008-06-24 12:03:19.000000000 -0700
@@ -9,4 +9,5 @@ obj_mq-$(CONFIG_COMPAT) += compat_mq.o
obj-$(CONFIG_POSIX_MQUEUE) += mqueue.o msgutil.o $(obj_mq-y)
obj-$(CONFIG_IPC_NS) += namespace.o
obj-$(CONFIG_MQ_NS) += mq_namespace.o
+obj-$(CONFIG_POSIX_MQUEUE_SYSCTL) += mq_sysctl.o

diff -puN /dev/null ipc/mq_sysctl.c
--- /dev/null 2007-04-11 11:48:27.000000000 -0700
+++ linux-2.6.git-dave/ipc/mq_sysctl.c 2008-06-24 12:03:19.000000000 -0700
@@ -0,0 +1,117 @@
+/*
+ * Copyright (C) 2007 IBM Corporation
+ *
+ * Author: Cedric Le Goater <clg@fr.ibm.com>
+ *
+ * This program is free software; you can redistribute it and/or
+ * modify it under the terms of the GNU General Public License as
+ * published by the Free Software Foundation, version 2 of the
+ * License.
+ */
+
+#include <linux/nsproxy.h>
+#include <linux/mq_namespace.h>
+#include <linux/sysctl.h>
+
+
+static void *get_mq(ctl_table *table)
+{
+    char *which = table->data;
+    struct mq_namespace *mq_ns = current->nsproxy->mq_ns;
+    which = (which - (char *)&init_mq_ns) + (char *)mq_ns;
+    return which;
+}
+
+ifdef CONFIG_PROC_FS
+static int proc_mq_dointvec(ctl_table *table, int write, struct file *filp,
+    void __user *buffer, size_t *lenp, loff_t *ppos)
+{

```

```

+ struct ctl_table mq_table;
+ memcpy(&mq_table, table, sizeof(mq_table));
+ mq_table.data = get_mq(table);
+
+ return proc_dointvec(&mq_table, write, filp, buffer, lenp, ppos);
+}
+
+static int proc_mq_dointvec_minmax(ctl_table *table, int write,
+ struct file *filp, void __user *buffer, size_t *lenp, loff_t *ppos)
+{
+ struct ctl_table mq_table;
+ memcpy(&mq_table, table, sizeof(mq_table));
+ mq_table.data = get_mq(table);
+
+ return proc_dointvec_minmax(&mq_table, write, filp, buffer,
+ lenp, ppos);
+}
+#else
+static int proc_mq_dointvec(ctl_table *table, int write, struct file *filp,
+ void __user *buffer, size_t *lenp, loff_t *ppos)
+{
+ return -ENOSYS;
+}
+
+static int proc_mq_dointvec_minmax(ctl_table *table, int write,
+ struct file *filp, void __user *buffer, size_t *lenp, loff_t *ppos)
+{
+ return -ENOSYS;
+}
#endif
+
+static int msg_max_limit_min = DFLT_MSGMAX;
+static int msg_max_limit_max = HARD_MSGMAX;
+
+static int msg_maxsize_limit_min = DFLT_MSGSIZEMAX;
+static int msg_maxsize_limit_max = INT_MAX;
+
+static ctl_table mq_sysctls[] = {
+ {
+ .procname = "queues_max",
+ .data = &init_mq_ns.queues_max,
+ . maxlen = sizeof(int),
+ .mode = 0644,
+ .proc_handler = proc_mq_dointvec,
+ },
+ {
+ .procname = "msg_max",
+ .data = &init_mq_ns.msg_max,
+

```

```

+ . maxlen = sizeof(int),
+ . mode = 0644,
+ . proc_handler = proc_mq_dointvec_minmax,
+ . extra1 = &msg_max_limit_min,
+ . extra2 = &msg_max_limit_max,
+ },
+ {
+ . procname = "msgsize_max",
+ . data = &init_mq_ns.msgsize_max,
+ . maxlen = sizeof(int),
+ . mode = 0644,
+ . proc_handler = proc_mq_dointvec_minmax,
+ . extra1 = &msg_maxsize_limit_min,
+ . extra2 = &msg_maxsize_limit_max,
+ },
+ { .ctl_name = 0 }
+};
+
+static ctl_table mq_sysctl_dir[] = {
+ {
+ .procname = "mqueue",
+ .mode = 0555,
+ .child = mq_sysctls,
+ },
+ { .ctl_name = 0 }
+};
+
+static ctl_table mq_sysctl_root[] = {
+ {
+ .ctl_name = CTL_FS,
+ .procname = "fs",
+ .mode = 0555,
+ .child = mq_sysctl_dir,
+ },
+ { .ctl_name = 0 }
+};
+
+struct ctl_table_header *mq_register_sysctl_table(void)
+{
+ return register_sysctl_table(mq_sysctl_root);
+}
diff -puN ipc/mqueue.c~mq_namespace-fix-sysctl ipc/mqueue.c
--- linux-2.6.git/ipc/mqueue.c~mq_namespace-fix-sysctl 2008-06-24 12:03:19.000000000 -0700
+++ linux-2.6.git-dave/ipc/mqueue.c 2008-06-24 12:03:19.000000000 -0700
@@ -1230,60 +1230,6 @@

@@ -1230,60 +1230,6 @@ struct file_system_type mqueue_fs_type =
    .kill_sb = kill_litter_super,
};


```

```

-static int msg_max_limit_min = DFLT_MSGMAX;
-static int msg_max_limit_max = HARD_MSGMAX;
-
-static int msg_maxsize_limit_min = DFLT_MSGSIZEMAX;
-static int msg_maxsize_limit_max = INT_MAX;
-
-static ctl_table mq_sysctls[] = {
-{
- .procname = "queues_max",
- .data = &init_mq_ns.queues_max,
- . maxlen = sizeof(int),
- .mode = 0644,
- .proc_handler = &proc_dointvec,
- },
- {
- .procname = "msg_max",
- .data = &init_mq_ns.msg_max,
- . maxlen = sizeof(int),
- .mode = 0644,
- .proc_handler = &proc_dointvec_minmax,
- .extra1 = &msg_max_limit_min,
- .extra2 = &msg_max_limit_max,
- },
- {
- .procname = "msgsize_max",
- .data = &init_mq_ns.msgsize_max,
- . maxlen = sizeof(int),
- .mode = 0644,
- .proc_handler = &proc_dointvec_minmax,
- .extra1 = &msg_maxsize_limit_min,
- .extra2 = &msg_maxsize_limit_max,
- },
- { .ctl_name = 0 }
-};

-
-static ctl_table mq_sysctl_dir[] = {
-{
- .procname = "mqueue",
- .mode = 0555,
- .child = mq_sysctls,
- },
- { .ctl_name = 0 }
-};

-
-static ctl_table mq_sysctl_root[] = {
-{
- .ctl_name = CTL_FS,
- .procname = "fs",

```

```
- .mode = 0555,
- .child = mq_sysctl_dir,
- },
- { .ctl_name = 0 }
-};
-
static int __init init_mqueue_fs(void)
{
    int error;
@@ -1295,7 +1241,7 @@ static int __init init_mqueue_fs(void)
    return -ENOMEM;

/* ignore failures - they are not fatal */
- mq_sysctl_table = register_sysctl_table(mq_sysctl_root);
+ mq_sysctl_table = mq_register_sysctl_table();

error = register_filesystem(&mqueue_fs_type);
if (error)
```

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
