

---

Subject: Re: [PATCH] mtd\_blkdevs: Convert to use the kthread API

Posted by [Christoph Hellwig](#) on Sun, 22 Apr 2007 12:24:53 GMT

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

---

On Thu, Apr 19, 2007 at 12:55:28AM -0600, Eric W. Biederman wrote:

> From: Eric W. Biederman <ebiederm@xmission.com> - unquoted

>

> thread\_run is used instead of kernel\_thread, daemonize, and mucking

> around blocking signals directly.

This is the full conversion I sent to Dave in April 2006, but never got any feedback to:

Signed-off-by: Christoph Hellwig <hch@lst.de>

Index: linux-2.6/drivers/mtd/mtd\_blkdevs.c

=====

--- linux-2.6.orig/drivers/mtd/mtd\_blkdevs.c 2007-01-29 10:03:52.000000000 +0100

+++ linux-2.6/drivers/mtd/mtd\_blkdevs.c 2007-04-22 13:22:03.000000000 +0200

@@ -20,6 +20,7 @@

#include <linux/hdreg.h>

#include <linux/init.h>

#include <linux/mutex.h>

+#include <linux/kthread.h>

#include <asm/uaccess.h>

static LIST\_HEAD(blktrans\_majors);

@@ -28,9 +29,7 @@ extern struct mutex mtd\_table\_mutex;

extern struct mtd\_info \*mtd\_table[];

struct mtd\_blkcore\_priv {

- struct completion thread\_dead;

- int exiting;

- wait\_queue\_head\_t thread\_wq;

+ struct task\_struct \*thread;

struct request\_queue \*rq;

spinlock\_t queue\_lock;

};

@@ -83,38 +82,19 @@ static int mtd\_blktrans\_thread(void \*arg

/\* we might get involved when memory gets low, so use PF\_MEMALLOC \*/

current->flags |= PF\_MEMALLOC | PF\_NOFREEZE;

- daemonize("%sd", tr->name);

-

- /\* daemonize() doesn't do this for us since some kernel threads

- actually want to deal with signals. We can't just call

- exit\_sighand() since that'll cause an oops when we finally

```

- do_exit. */
- spin_lock_irq(&current->sigband->siglock);
- sigfillset(&current->blocked);
- recalc_sigpending();
- spin_unlock_irq(&current->sigband->siglock);
-
  spin_lock_irq(rq->queue_lock);
-
- while (!tr->blkcore_priv->exiting) {
+ while (!kthread_should_stop()) {
  struct request *req;
  struct mtd_blktrans_dev *dev;
  int res = 0;
- DECLARE_WAITQUEUE(wait, current);

  req = elv_next_request(rq);

  if (!req) {
- add_wait_queue(&tr->blkcore_priv->thread_wq, &wait);
- set_current_state(TASK_INTERRUPTIBLE);
-
  spin_unlock_irq(rq->queue_lock);
-
+ set_current_state(TASK_INTERRUPTIBLE);
  schedule();
- remove_wait_queue(&tr->blkcore_priv->thread_wq, &wait);
-
  spin_lock_irq(rq->queue_lock);
-
  continue;
  }

@@ -133,13 +113,13 @@ static int mtd_blktrans_thread(void *arg
  }
  spin_unlock_irq(rq->queue_lock);

- complete_and_exit(&tr->blkcore_priv->thread_dead, 0);
+ return 0;
  }

static void mtd_blktrans_request(struct request_queue *rq)
{
  struct mtd_blktrans_ops *tr = rq->queuedata;
- wake_up(&tr->blkcore_priv->thread_wq);
+ wake_up_process(tr->blkcore_priv->thread);
  }

```

```

@@ -388,8 +368,6 @@ int register_mtd_blktrans(struct mtd_blk
    return ret;
}
spin_lock_init(&tr->blkcore_priv->queue_lock);
- init_completion(&tr->blkcore_priv->thread_dead);
- init_waitqueue_head(&tr->blkcore_priv->thread_wq);

tr->blkcore_priv->rq = blk_init_queue(mtd_blktrans_request, &tr->blkcore_priv->queue_lock);
if (!tr->blkcore_priv->rq) {
@@ -403,13 +381,14 @@ int register_mtd_blktrans(struct mtd_blk
    blk_queue_hardsect_size(tr->blkcore_priv->rq, tr->blksize);
    tr->blkshift = ffs(tr->blksize) - 1;

- ret = kernel_thread(mtd_blktrans_thread, tr, CLONE_KERNEL);
- if (ret < 0) {
+ tr->blkcore_priv->thread = kthread_run(mtd_blktrans_thread, tr,
+ "%sd", tr->name);
+ if (IS_ERR(tr->blkcore_priv->thread)) {
    blk_cleanup_queue(tr->blkcore_priv->rq);
    unregister_blkdev(tr->major, tr->name);
    kfree(tr->blkcore_priv);
    mutex_unlock(&mtd_table_mutex);
- return ret;
+ return PTR_ERR(tr->blkcore_priv->thread);
}

INIT_LIST_HEAD(&tr->devs);
@@ -432,9 +411,7 @@ int deregister_mtd_blktrans(struct mtd_b
    mutex_lock(&mtd_table_mutex);

/* Clean up the kernel thread */
- tr->blkcore_priv->exiting = 1;
- wake_up(&tr->blkcore_priv->thread_wq);
- wait_for_completion(&tr->blkcore_priv->thread_dead);
+ kthread_stop(tr->blkcore_priv->thread);

/* Remove it from the list of active majors */
list_del(&tr->list);

```

---

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

---