

[PATCH 1/1] block: CFQ refcounting fix

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-08/7741.html>

brking_at_us.ibm.com

Date: 08/31/05

To: axboe@suse.de

Date: Tue, 30 Aug 2005 17:41:07 -0500

I ran across a memory leak related to the cfq scheduler. The cfq init function increments the refcnt of the associated request_queue. This refcount gets decremented in cfq's exit function. Since blk_cleanup_queue only calls the elevator exit function when its refcnt goes to zero, the request_q never gets cleaned up. It didn't look like other io schedulers were incrementing this refcnt, so I removed the refcnt increment and it fixed the memory leak for me.

To reproduce the problem, simply use cfq and use the scsi_host scan sysfs attribute to scan " - - - " repeatedly on a scsi host and watch the memory vanish.

Signed-off-by: Brian King <brking@us.ibm.com>

```
linux-2.6-bjking1/drivers/block/cfq-iosched.c | 1 -
1 files changed, 1 deletion(-)
diff -puN drivers/block/cfq-iosched.c~cfq_refcnt_fix drivers/block/cfq-iosched.c
--- linux-2.6/drivers/block/cfq-iosched.c~cfq_refcnt_fix      2005-08-30 17:26:55.000000000 -0500
+++ linux-2.6-bjking1/drivers/block/cfq-iosched.c      2005-08-30 17:26:55.000000000 -0500
@@ -2318,7 +2318,6 @@ static int cfq_init_queue(request_queue_
     e->elevator_data = cfqd;

     cfqd->queue = q;
-    atomic_inc(&q->refcnt);

     cfqd->max_queued = q->nr_requests / 4;
     q->nr_batching = cfq_queued;
```

-

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>