

---

Subject: [RFC][PATCH 7/9] cgroups: block: cfq: I/O bandwidth controlling subsystem for CGroups based on CFQ

Posted by [Vasily Tarasov](#) on Fri, 15 Feb 2008 06:59:53 GMT

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

---

From: Vasily Tarasov <[vtaras@openvz.org](mailto:vtaras@openvz.org)>

Takes into account requests that are in driver now. Before that we switched cgroup if there were no requests to service in this cgroup. Now we also check that there are no requests from this cgroups currently processed by driver and only in this case allow switching.

Signed-off-by: Vasily Tarasov <[vtaras@openvz.org](mailto:vtaras@openvz.org)>

---

```
--- linux-2.6.25-rc5-mm1/include/linux/cfq-iosched.h.rqindrv 2008-02-15 01:10:42.000000000 -0500
+++ linux-2.6.25-rc5-mm1/include/linux/cfq-iosched.h 2008-02-15 01:11:45.000000000 -0500
@@ -57,6 +57,7 @@ struct cfqio_cgroup_data {
 /* rr list of queues with requests */
 struct cfq_rb_root service_tree;
 unsigned long rqnum;
+unsigned long on_dispatch;
};

/*
--- linux-2.6.25-rc5-mm1/block/cfq-iosched.c.rqindrv 2008-02-15 01:10:42.000000000 -0500
+++ linux-2.6.25-rc5-mm1/block/cfq-iosched.c 2008-02-15 01:11:45.000000000 -0500
@@ -823,6 +823,7 @@ static void cfq_dispatch_insert(struct r
 cfq_remove_request(rq);
 cfqq->dispatched++;
+cfqq->cfqio_cgrp->on_dispatch++;
 elv_dispatch_sort(q, rq);

 if (cfq_cfqq_sync(cfqq))
@@ -1775,6 +1776,7 @@ static void cfq_completed_request(struct
 WARN_ON(!cfqq->dispatched);
 cfqd->rq_in_driver--;
 cfqq->dispatched--;
+cfqq->cfqio_cgrp->on_dispatch--;

 if (cfq_cfqq_sync(cfqq))
 cfqd->sync_flight--;
--- linux-2.6.25-rc5-mm1/block/cfqio-cgroup.c.rqindrv 2008-02-15 01:10:42.000000000 -0500
+++ linux-2.6.25-rc5-mm1/block/cfqio-cgroup.c 2008-02-15 01:11:45.000000000 -0500
@@ -54,7 +54,8 @@ static inline void set_active_cgrp(struc
```

```
void cfqio_cgrp_schedule_active(struct cfq_data *cfqd)
{
    if (cfqio_cgrp_expired(cfqd) || !cfqd->active_cfqio_cgroup ||
-    !cfqd->active_cfqio_cgroup->rqnum)
+ (!cfqd->active_cfqio_cgroup->rqnum &&
+ !cfqd->active_cfqio_cgroup->on_dispatch))
    set_active_cgrp(cfqd);
}
```

---