

---

Subject: [PATCH] block layer: ioprio\_best function fix

Posted by [Anonymous Coward](#) on Thu, 12 Oct 2006 12:13:30 GMT

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

---

Currently ioprio\_best function first checks whether aioprio or bioprio equals IOPRIO\_CLASS\_NONE (ioprio\_valid() macro does that) and if it is so it returns bioprio/aioprio appropriately. Thus the next four lines, that set aclass/bclass to IOPRIO\_CLASS\_BE, if aclass/bclass == IOPRIO\_CLASS\_NONE, are never executed.

The second problem: if aioprio from class IOPRIO\_CLASS\_NONE and bioprio from class IOPRIO\_CLASS\_IDLE are passed to ioprio\_best function, it will return IOPRIO\_CLASS\_IDLE. It means that during \_\_make\_request we can merge two requests and set the priority of merged request to IDLE, while one of the initial requests originates from a process with NONE (default) priority. So we can get a situation when a process with default ioprio will experience IO starvation, while there is no process from real-time class in the system.

Just removing ioprio\_valid check should correct situation.

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

--

```
--- linux-2.6.18/fs/ioprio.c.orig 2006-09-20 07:42:06.000000000 +0400
+++ linux-2.6.18/fs/ioprio.c 2006-10-12 14:52:36.000000000 +0400
@@ -145,11 +145,6 @@ int ioprio_best(unsigned short aprio, un
     unsigned short aclass = IOPRIO_PRIO_CLASS(aprio);
     unsigned short bclass = IOPRIO_PRIO_CLASS(bprio);

- if (!ioprio_valid(aprio))
-     return bprio;
- if (!ioprio_valid(bprio))
-     return aprio;
-
+ if (aclass == IOPRIO_CLASS_NONE)
+     aclass = IOPRIO_CLASS_BE;
+ if (bclass == IOPRIO_CLASS_NONE)
```

---