
Subject: [RFC][PATCH 0/9] cgroups: block: cfq: I/O bandwidth controlling subsystem for CGroups based on CFQ

Posted by [Vasily Tarasov](#) on Fri, 15 Feb 2008 06:53:34 GMT

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

From: Vasily Tarasov <vtaras@openvz.org>

The following patchset introduces I/O bandwidth controlling subsystem for the CGroups framework based on the CFQ scheduler.

User can assign a priority from 0 to 7 to a cgroup, and I/O bandwidth will be adjusted proportionally.

This effect is achieved by introducing a certain modifications to CFQ.

Now there is not only per process time slice, but also per cgroup time slice. During cgroup's time slice only processes from current cgroup can add requests to the queue. Inside cgroups's time slice processes are managed based on (per-process) time slices (as usual in CFQ). This is why we call this approach "two level CFQ".

Such kind of system works in OpenVZ project.

Patch is against 2.6.25-rc2-mm1

Comments, suggestions, criticism are all welcome.
