
Subject: [PATCH -mm 0/3] cgroup: block device i/o bandwidth controller (v5)

Posted by [Andrea Righi](#) on Sat, 12 Jul 2008 11:31:28 GMT

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

The objective of the i/o bandwidth controller is to improve i/o performance predictability of different cgroups sharing the same block devices.

Respect to other priority/weight-based solutions the approach used by this controller is to explicitly choke applications' requests that directly (or indirectly) generate i/o activity in the system.

The direct bandwidth limiting method has the advantage of improving the performance predictability at the cost of reducing, in general, the overall performance of the system (in terms of throughput).

Detailed informations about design, its goal and usage are described in the documentation.

Tested against 2.6.26-rc8-mm1.

The all-in-one patch (and previous versions) can be found at:
<http://download.systemimager.org/~arighi/linux/patches/io-throttle/>

Thanks to Li Zefan for reviewing.

Changelog: (v4 -> v5)

- rebase to 2.6.26-rc8-mm1
- handle AIO opportunely: return -EAGAIN from sys_io_submit(), instead of making to sleep tasks doing AIO
- userspace=>kernel interface now accepts the following syntaxes:
 - * dev:0 <- to delete a limiting rule
 - * dev:bw-limit:0 <- define a leaky bucket throttling rule
 - * dev:bw-limit:1:bucket-size <- define a token bucket throttling rule
- use .write_string and .read_seq_string to simplify iothrottle_read() and iothrottle_write() functions
- use a enum structure to enumerate the various throttling algorithms

TODO:

- see documentation

-Andrea

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
