
Subject: Re: [PATCH 2/2] dm-ioband: I/O bandwidth controller v0.0.4: Document
Posted by [akpm](#) on Sat, 26 Apr 2008 18:45:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 24 Apr 2008 20:22:19 +0900 (JST) Ryo Tsuruta <ryov@valinux.co.jp> wrote:

> +What's dm-ioband all about?

> +

> + dm-ioband is an I/O bandwidth controller implemented as a device-mapper
> + driver. Several jobs using the same physical device have to share the
> + bandwidth of the device. dm-ioband gives bandwidth to each job according
> + to its weight, which each job can set its own value to.

> +

> + At this time, a job is a group of processes with the same pid or pgrp or
> + uid. There is also a plan to make it support cgroup. A job can also be a
> + virtual machine such as KVM or Xen.

Most writes are performed by pdflush, kswapd, etc. This will lead to large inaccuracy.

It isn't trivial to fix. We'd need deep, long tracking of ownership probably all the way up to the pagecache page. The same infrastructure would be needed to make Sergey's "BSD acct: disk I/O accounting" vaguely accurate. Other proposals need it, but I forget what they are.

Much more minor points: when merge-time comes, the patches should have the LINUX_VERSION_CODE stuff removed. And probably all of the many `inline's should be removed.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
