Subject: Re: [RFC][v2][patch 0/12][CFQ-cgroup]Yet another I/O bandwidth controlling subsystem for CGroups bas
Posted by Ryo Tsuruta on Fri, 25 Apr 2008 09:54:44 GMT
View Forum Message <> Reply to Message

Hi,

I report benchmark results of the following I/O bandwidth controllers.

> From: Vasily Tarasov <vtaras@openvz.org>
> Subject: [RFC][PATCH 0/9] cgroups: block: cfq: I/O bandwidth
> controlling subsystem for CGroups based on CFQ
> Date: Fri, 15 Feb 2008 01:53:34 -0500
>
> From: "Satoshi UCHIDA" <s-uchida@ap.jp.nec.com>
> Subject: [RFC][v2][patch 0/12][CFQ-cgroup]Yet another I/O bandwidth
> controlling subsystem for CGroups based on CFQ
> Date: Thu, 3 Apr 2008 16:09:12 +0900

The test procedure is as follows:
- o Prepare 3 partitions sdc2, sdc3 and sdc4.
- o Run 100 processes issuing random direct I/O with 4KB data on each partitions.
- o Run 3 tests:
  - #1 issuing read I/O only.
  - #2 issuing write I/O only.
  - #3 sdc2 and sdc3 are read, sdc4 is write.
- o Count up the number of I/Os which have done in 60 seconds.

Unfortunately, both bandwidth controllers didn't work as I expected,
On the test #3, the write I/O ate up the bandwidth regardless of the
specified priority level.

Vasily's scheduler
The number of I/Os (percentage to total I/Os)

partition	sdc2	sdc3	sdc4	total
priority	7(highest)	4	0(lowest)	I/Os
#1 read	3620(35.6%)	3474(34.2%)	3065(30.2%)	10159
#2 write	21985(36.6%)	19274(32.1%)	18856(31.4%)	60115
#3 read&write	5571(7.5%)	3253(4.4%)	64977(88.0%)	73801

Satoshi's scheduler
The number of I/Os (percentage to total I/O)

partition	sdc2	sdc3	sdc4	total

priority	0(highest)	4	7(lowest)	I/Os
#1 read	4523(47.8%)	3733(39.5%)	1204(12.7%)	9460
#2 write	65202(59.0%)	35603(32.2%)	9673(8.8%)	110478
#3 read&write	5328(23.0%)	4153(17.9%)	13694(59.1%)	23175

I'd like to see other benchmark results if anyone has.

Thanks,
Ryo Tsuruta