
Subject: Re: Re: [RFC][v2][patch 0/12][CFQ-cgroup]Yet another I/O bandwidth controlling subsystem for CGroup

Posted by [Florian Westphal](#) on Fri, 25 Apr 2008 21:37:02 GMT

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

Ryo Tsuruta <ryov@valinux.co.jp> wrote:

[..]

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

Here are a few results. IO is issued in 4k chunks, using O_DIRECT. Each process issues both reads and writes. There are 60 such processes in each cgroup (except where noted). Numbers given show the total count of io requests (read and write) completed in 60 seconds. All processes use the same partition, fs is ext3.

Vasily's scheduler:

```
-----
| cgroup | s0          | s1          | total |
| priority | 4          | 4          | |I/Os |
-----
|      | 24953      | 24062      | 49015|
|      | 29558(60 processes)| 14639 (30 proc)| 44197|
-----
| priority | 0          | 4          |      |
|      | 24221      | 24047      | 48268|
| priority | 1          | 4          |      |
|      | 24897      | 24509      | 49406|
| priority | 2          | 4          |      |
|      | 23295      | 23622      | 46917|
| priority | 0          | 7          |      |
|      | 22301      | 23373      | 45674|
-----
```

Satoshi's scheduler:

```
-----
| cgroup | s0          | s1          | total |
| priority | 3          | 3          | |I/Os |
-----
|      | 25175      | 26463      | 51638|
|      | 26944 (60) | 26698 (30) | 53642|
-----
| priority | 0          | 3          |      |
|      | 60821      | 19846      | 80667|
| priority | 1          | 3          |      |
|      | 50608      | 25994      | 76602|
| priority | 2          | 3          |      |
|      | 32132      | 26641      | 58773|
| priority | 7          | 0          |      |
-----
```

| | 91387 | 12547 |103934|

So in short, i can't see any effect when i use Vasily's i/o scheduler. Setting
echo 10 > /sys/block/hda/queue/iosched/cgrp_slice
did at least show different results in the 'prio 7 vs. prio 0 case' (~29000 (prio 7) vs. 20000 (prio 0)).

What i found surprising is that Satoshis scheduler has about twice of the io count...

Thanks, Florian
