

---

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

Posted by [Ryo Tsuruta](#) on Tue, 29 Apr 2008 00:44:38 GMT

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

---

Hi Florian,

> Here are a few results.

Thank you very much.

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

Of my previous results, Satoshi's scheduler has about twice of the io count on the write test, but both io counts on the read test are nearly the same.

Here are another results. The test procedure is as follows:

- o Prepare 3 partitions sdc2, sdc3 and sdc4.
- o Run 50 processes for sdc2, 100 processes for sdc3 and 150 processes for sdc4. Apply I/O loads in inverse to the priority level.
- o Each process issuing random read/write direct I/O with 4KB data.
- o Count up the number of I/Os which have done in 60 seconds.

The number of I/Os (percentage to total I/Os)

|                               |            |           |           |       |
|-------------------------------|------------|-----------|-----------|-------|
| -----                         |            |           |           |       |
| partition                     | sdc2       | sdc3      | sdc4      | total |
| processes                     | 50         | 100       | 150       | I/Os  |
| -----+-----+-----+-----+----- |            |           |           |       |
| Normal                        | 2281(18%)  | 4287(34%) | 6005(48%) | 12573 |
| -----+-----+-----+-----+----- |            |           |           |       |
| Vasily's sched.               |            |           |           |       |
| cgroup priority               | 7(highest) | 4         | 0(lowest) |       |
|                               | 3713(24%)  | 6587(42%) | 5354(34%) | 15654 |
| -----+-----+-----+-----+----- |            |           |           |       |
| Satoshi's sched.              |            |           |           |       |
| cgroup priority               | 0(highest) | 4         | 7(lowest) |       |
|                               | 5399(42%)  | 6506(50%) | 1034( 8%) | 12939 |
| -----                         |            |           |           |       |

Satoshi's scheduler suppressed the I/O to the lowest priority partition better than Vasily's one, but Vasily's scheduler got the highest total I/Os.

Thanks,  
Ryo Tsuruta

---