

## Re: [PATCH] O11int for interactivity

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-08/1022.html>

---

**From:** Andrew Morton ([akpm\\_at\\_osdl.org](mailto:akpm_at_osdl.org))

**Date:** 08/05/03

Date: Mon, 4 Aug 2003 22:55:32 -0700

To: [Valdis.Kletnieks@vt.edu](mailto:Valdis.Kletnieks@vt.edu)

[Valdis.Kletnieks@vt.edu](mailto:Valdis.Kletnieks@vt.edu) wrote:

>

> *The \*odd\* part is that the pgpgin, pgpgout, and pswpin numbers do \*NOT\* seem to be correlated. High I/O loads from read/write don't seem to cause a problem – untarring the Linux distro won't do it, running badblocks won't do it.*

>

> *But if somebody has to swap out, all hell breaks loose...*

swapout tends to happen via page reclaim, whereas normal writeback does not.

What's the difference? When swapout is happening you can expect increased latency in the page allocator.

My guess is that xmms is getting throttled in `try_to_free_pages()`.

There is a very good argument for giving `!SCHED_OTHER` tasks "special treatment" in the VM. ie:

a) exempt them from `balance_dirty_pages()` throttling treatment altogether

b) let them dip further into the page reserves in `__alloc_pages`.

iirc, -aa kernels do some of this. As does the Digeo kernel. Just haven't got around to it in 2.6. It's pretty simple.

If xmms isn't running `SCHED_FIFO/SCHED_RR`, well, you lose.

The instrumentation to add is page allocation latency.

Another possibility is that xmms is getting stuck in a read. The anticipatory scheduler is currently rather tuned for throughput. Judging by the vmstat trace which was posted, we have a classic read-stream-vs-write-stream going on. We trade off latency versus throughput; perhaps wrongly. You can decrease latency (at the expense of throughput) by decreasing the settings in `/sys/block/hda/queue/iosched`.

Linux-Kernel: Re: [PATCH] O11int for interactivity

To a point, it is a nice linear tradeoff, and someone should put the time in to tweak and characterise it.

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to [majordomo@vger.kernel.org](mailto:majordomo@vger.kernel.org)

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>