

Re: Make pipe data structure be a circular list of pages, rather than

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-01/2377.html>

From: Lee Revell (rlrevell_at_joe-job.com)

Date: 01/08/05

To: Chris Friesen <cfriesen@nortelnetworks.com>

Date: Sat, 08 Jan 2005 16:38:45 -0500

On Fri, 2005-01-07 at 17:46 -0600, Chris Friesen wrote:

> *Mike Waychison wrote:*

>

> > *This got me to thinking about how you can heuristically optimize away*
> > *coalescing support and still allow PAGE_SIZE bytes minimum in the*
> > *effective buffer.*

>

> *While coalescing may be a win in some cases, there should also be some*
> *way to tell the kernel to NOT coalesce, to handle the case where you*
> *want minimum latency at the cost of some throughput.*

Many latency critical apps use (tmpfs mounted) FIFO's for IPC; the Linux FIFO being one of the fastest known IPC mechanisms. Each client in the JACK (<http://jackit.sf.net>) graph wakes the next one by writing a single byte to a FIFO. Ardour's GUI, control, and audio threads interact via a similar mechanism. How would you expect this change to impact the inter thread wakeup latency? It's confusing when people say "performance", meaning "increased throughput albeit with more latency". For many people that's a regression.

These apps **certainly** care about performance, they just don't define it in terms of throughput.

And yes we do know futexes are the right tool for this but they weren't available at the time and aren't available on 2.4.

Lee

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to 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/>