

Re: [RFC] kernel facilities for cache prefetching

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-05/msg00384.html>

- *From:* Wu Fengguang <wfg@xxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Tue, 2 May 2006 16:53:25 +0800
-

On Tue, May 02, 2006 at 10:30:17AM +0200, Arjan van de Ven wrote:

one interesting thing that came out of the fedora readahead work is that most of the bootup isn't actually IO bound. My test machine for example can load all the data into ram in about 10 seconds, so that the rest of the boot is basically IO-less. But that still takes 2 minutes....
So I'm not entirely sure how much you can win by just attacking this.

Yes, I find it hard to improve the boot time of the init.d stage. However, it is perfectly ok to preload all GUI stuffs during that timespan, by overlapping CPU/IO activities.

Another interesting approach would be to actually put all the data you want to use in a non-fragmented, sequential area on disk somehow (there is an OLS paper submitted about that by Ben) so that at least the disk side is seekless...

You are right, reducing seeking distances helps not much. My fluxbox desktop requires near 3k seeks, which can be loaded in the 20s init.d booting time. But for KDE/GNOME desktops, some defragging would be necessary to fit them into the 20s time span.

I found ext3 to be rather good filesystem to support poor man's defrag described by Chris Mason:

<http://www.gelato.unsw.edu.au/archives/git/0504/1690.html>

This certainly can help. Based on some ideas from andrea I made a poor man's defrag script last year that was similar. It worked by copying files into a flat dir in the order you expected to read them in, deleting the original, then hard linking them into their original name.

Make the 'flat dir' as an top level dir would do the trick for ext3. We have good chance to merge 10k seeks into 3k seeks by this trick :)

Wu

—

Re: [RFC] kernel facilities for cache prefetching

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>