

Re: [Ext2-devel] [RFC] [PATCH] Reducing average ext2 fsck time through fs-wide dirty bit]

Re: [Ext2-devel] [RFC] [PATCH] Reducing average ext2 fsck time through fs-wide dirty bit]

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-03/msg08632.html>

- *From:* Andreas Dilger <adilger@xxxxxxxxxxxxxx>
 - *Date:* Fri, 24 Mar 2006 12:31:25 -0700
-

On Mar 24, 2006 11:13 -0800, Mingming Cao wrote:

There are reasons for zeroing indirect blocks on truncate:

- * There are limits to the size of a single journal transaction (1/4 of the journal size). When truncating a large fragmented file, it may require modifying so many block bitmaps and group descriptors that it forces a journal transaction to close out, stalling the unlink operation.
- * Because of this per-transaction limit, truncate needs to zero the [dt]indirect blocks starting from the end of the file, in case it needs to start a new transaction in the middle of the truncate (ext3 guarantees that a partially-completed truncate will be consistent/completed after a crash).
- * The read/write of the file's [dt]indirect blocks from the end of the file to the beginning can take a lot of time, as it does this in single-block chunks and the blocks are not contiguous.

See my recent post on how this performance problem could be fixed.

Cheers, Andreas

Andreas Dilger
Principal Software Engineer
Cluster File Systems, Inc.

—

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