

device mapper not reporting no-barrier-support?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-02/msg12842.html>

- *From:* Anders Henke <anders.henke@xxxxxxxx>
 - *Date:* Mon, 25 Feb 2008 14:26:15 +0100
-

Hi,

I'm currently stuck between Kernel LVM and DRBD, as I'm using Kernel 2.6.24.2 with DRBD 8.2.5 on top of an LVM2 device (LV).

-LVM2/device mapper doesn't support write barriers
-DRBD uses `blkdev_issue_flush()` to flush its metadata to disk.
On a no-barrier-device, DRBD should receive EOPNOTSUPP, but it really does receive an EIO. Promptly, DRBD gives the error message "drbd0: local disk flush failed with status -5".

The physical disk (in LVM speak) is a RAID1 on a 3ware 9650SE-2LP controller; the driver 3w-9xxx supports barriers and after moving my DRBD device from the LV to a single partition on the same RAID1, the error messages from DRBD vanished.

I've posted a lengthy summary of my findings to

<http://lists.linbit.com/pipermail/drbd-user/2008-February/008665.html>

... where Lars Ellenberg from DRBD basically responded in

<http://lists.linbit.com/pipermail/drbd-user/2008-February/008666.html>

... that DRBD does catch the EOPNOTSUPP for `blkdev_issue_flush` and `BIO_RW_BARRIER`, but the lvm implementation of `blkdev_issue_flush` in 2.6.24.2 apparently does return EIO for `blkdev_issue_flush`.

So simply the question: how should a top-layer driver check whether a lower device does support barriers? `md-raid` does check this way differently than e.g. XFS does, while DRBD also adds a third way to check this.
Or is this "merely" a bug in `drivers/md/dm.c`?

Anders

1&1 Internet AG System Architect
Braucherstrasse 48 v://49.721.91374.50
D-76135 Karlsruhe f://49.721.91374.225

device mapper not reporting no-barrier-support?

Amtsgericht Montabaur HRB 6484

Vorstand: Henning Ahlert, Ralph Dommermuth, Matthias Ehrlich, Andreas Gauger,
Thomas Gottschlich, Matthias Greve, Robert Hoffmann, Markus Huhn, Achim Weiss
Aufsichtsratsvorsitzender: Michael Scheeren

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
the body of a message to majordomo@xxxxxxxxxxxxxxxxxx

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

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