

```

*****
44738 Wed Nov  6 12:52:06 2013
new/usr/src/uts/common/fs/zfs/dmu_tx.c
4188 assertion failed in dmu_tx_hold_free(): dn_datablkshift != 0
Reviewed by: George Wilson <george.wilson@delphix.com>
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
*****
unchanged_portion_omitted_

586 void
587 dmu_tx_hold_free(dmu_tx_t *tx, uint64_t object, uint64_t off, uint64_t len)
588 {
589     dmu_tx_hold_t *txh;
590     dnode_t *dn;
591     int err;
592     zio_t *zio;

594     ASSERT(tx->tx_txg == 0);

596     txh = dmu_tx_hold_object_impl(tx, tx->tx_objset,
597     object, THT_FREE, off, len);
598     if (txh == NULL)
599         return;
600     dn = txh->txh_dnode;
601     dmu_tx_count_dnode(txh);

603     if (off >= (dn->dn_maxblkid+1) * dn->dn_datablksz)
604         return;
605     if (len == DMU_OBJECT_END)
606         len = (dn->dn_maxblkid+1) * dn->dn_datablksz - off;

608     /*
609     * For i/o error checking, we read the first and last level-0
610     * blocks if they are not aligned, and all the level-1 blocks.
611     *
612     * Note: dbuf_free_range() assumes that we have not instantiated
613     * any level-0 dbufs that will be completely freed. Therefore we must
614     * exercise care to not read or count the first and last blocks
615     * if they are blocksize-aligned.
616     */
617     if (dn->dn_datablkshift == 0) {
618         if (off != 0 || len < dn->dn_datablksz)
619             dmu_tx_count_write(txh, 0, dn->dn_datablksz);
620     } else {
621         /* first block will be modified if it is not aligned */
622         if (!IS_P2ALIGNED(off, 1 << dn->dn_datablkshift))
623             dmu_tx_count_write(txh, off, 1);
624         /* last block will be modified if it is not aligned */
625         if (!IS_P2ALIGNED(off + len, 1 << dn->dn_datablkshift))
626             dmu_tx_count_write(txh, off+len, 1);
627     }

629     /*
630     * Check level-1 blocks.
631     */
632     if (dn->dn_nlevels > 1) {
633         int shift = dn->dn_datablkshift + dn->dn_indblkshift -
634             SPA_BLKPTRSHIFT;
635         uint64_t start = off >> shift;
636         uint64_t end = (off + len) >> shift;

638         ASSERT(dn->dn_datablkshift != 0);
638         ASSERT(dn->dn_indblkshift != 0);

640         /*
641         * dnode_reallocate() can result in an object with indirect

```

```

642     * blocks having an odd data block size. In this case,
643     * just check the single block.
644     */
645     if (dn->dn_datablkshift == 0)
646         start = end = 0;

648     zio = zio_root(tx->tx_pool->dp_spa,
649         NULL, NULL, ZIO_FLAG_CANFAIL);
650     for (uint64_t i = start; i <= end; i++) {
651         uint64_t ibyte = i << shift;
652         err = dnode_next_offset(dn, 0, &ibyte, 2, 1, 0);
653         i = ibyte >> shift;
654         if (err == ESRCH)
655             break;
656         if (err) {
657             tx->tx_err = err;
658             return;
659         }

661         err = dmu_tx_check_ioerr(zio, dn, 1, i);
662         if (err) {
663             tx->tx_err = err;
664             return;
665         }
666     }
667     err = zio_wait(zio);
668     if (err) {
669         tx->tx_err = err;
670         return;
671     }
672 }

674     dmu_tx_count_free(txh, off, len);
675 }
unchanged_portion_omitted_

```