

new/usr/src/lib/libzfs/common/libzfs\_dataset.c

1

```
*****
111785 Mon Aug 12 01:16:02 2013
new/usr/src/lib/libzfs/common/libzfs_dataset.c
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
_____unchanged_portion_omitted_____

3509 /*
3510 * Given a dataset, rollback to a specific snapshot, discarding any
3511 * data changes since then and making it the active dataset.
3512 *
3513 * Any snapshots more recent than the target are destroyed, along with
3514 * their dependents.
3515 */
3516 int
3517 zfs_rollback(zfs_handle_t *zhp, zfs_handle_t *snap, boolean_t force)
3518 {
3519     rollback_data_t cb = { 0 };
3520     int err;
3521     zfs_cmd_t zc = { 0 };
3522     boolean_t restore_resv = 0;
3523     uint64_t old_volsize, new_volsize;
3524     zfs_prop_t resv_prop;

3525     assert(zhp->zfs_type == ZFS_TYPE_FILESYSTEM ||
3526            zhp->zfs_type == ZFS_TYPE_VOLUME);

3528     /*
3529      * Destroy all recent snapshots and their dependents.
3530      */
3531     cb.cb_force = force;
3532     cb.cb_target = snap->zfs_name;
3533     cb.cb_create = zfs_prop_get_int(snap, ZFS_PROP_CREATETXG);
3534     (void) zfs_iter_children(zhp, rollback_destroy, &cb);

3536     if (cb.cb_error)
3537         return (-1);

3539     /*
3540      * Now that we have verified that the snapshot is the latest,
3541      * rollback to the given snapshot.
3542      */

3544     if (zhp->zfs_type == ZFS_TYPE_VOLUME) {
3545         if (zfs_which_resv_prop(zhp, &resv_prop) < 0)
3546             return (-1);
3547         old_volsize = zfs_prop_get_int(zhp, ZFS_PROP_VOLSIZE);
3548         restore_resv =
3549             (old_volsize == zfs_prop_get_int(zhp, resv_prop));
3550     }

3553     (void) strncpy(zc.zc_name, zhp->zfs_name, sizeof (zc.zc_name));

3555     if (ZFS_IS_VOLUME(zhp))
3556         zc.zc_objset_type = DMU_OST_ZVOL;
3557     else
3558         zc.zc_objset_type = DMU_OST_ZFS;

3552     /*
3553      * We rely on zfs_iter_children() to verify that there are no
3554      * newer snapshots for the given dataset. Therefore, we can
3555      * simply pass the name on to the ioctl() call. There is still
3556      * an unlikely race condition where the user has taken a
```

new/usr/src/lib/libzfs/common/libzfs\_dataset.c

2

```
3557     * snapshot since we verified that this was the most recent.
3558     */
3559     err = lzcb_rollback(zhp->zfs_name, NULL, 0);
3560     if (err != 0) {
3561         if ((err = zfs_ioctl(zhp->zfs_hdl, ZFS_IOC_ROLLBACK, &zc)) != 0) {
3562             (void) zfs_standard_error_fmt(zhp->zfs_hdl, errno,
3563             dgettext(TEXT_DOMAIN, "cannot rollback '%s'",
3564             zhp->zfs_name);
3565             return (err);
3566         }
3567     }
3568     /*
3569      * For volumes, if the pre-rollback volsize matched the pre-
3570      * rollback reservation and the volsize has changed then set
3571      * the reservation property to the post-rollback volsize.
3572      * Make a new handle since the rollback closed the dataset.
3573      */
3574     if ((zhp->zfs_type == ZFS_TYPE_VOLUME) &&
3575         (zhp = make_dataset_handle(zhp->zfs_hdl, zhp->zfs_name))) {
3576         if (restore_resv) {
3577             new_volsize = zfs_prop_get_int(zhp, ZFS_PROP_VOLSIZE);
3578             if (old_volsize != new_volsize)
3579                 err = zfs_prop_set_int(zhp, resv_prop,
3580                 new_volsize);
3581         }
3582         zfs_close(zhp);
3583     }
3584     return (err);
3585 }
_____unchanged_portion_omitted_____
```

new/usr/src/lib/libzfs\_core/common/libzfs\_core.c

1

```
*****
17237 Mon Aug 12 01:16:03 2013
new/usr/src/lib/libzfs_core/common/libzfs_core.c
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */
21
22 /*
23  * Copyright (c) 2013 by Delphix. All rights reserved.
24  * Copyright (c) 2012 by Delphix. All rights reserved.
25  * Copyright (c) 2013 Steven Hartland. All rights reserved.
26 */
27 /*
28  * LibZFS_Core (lzc) is intended to replace most functionality in libzfs.
29  * It has the following characteristics:
30  *
31  * - Thread Safe. libzfs_core is accessible concurrently from multiple
32  * threads. This is accomplished primarily by avoiding global data
33  * (e.g. caching). Since it's thread-safe, there is no reason for a
34  * process to have multiple libzfs "instances". Therefore, we store
35  * our few pieces of data (e.g. the file descriptor) in global
36  * variables. The fd is reference-counted so that the libzfs_core
37  * library can be "initialized" multiple times (e.g. by different
38  * consumers within the same process).
39  *
40  * - Committed Interface. The libzfs_core interface will be committed,
41  * therefore consumers can compile against it and be confident that
42  * their code will continue to work on future releases of this code.
43  * Currently, the interface is Evolving (not Committed), but we intend
44  * to commit to it once it is more complete and we determine that it
45  * meets the needs of all consumers.
46  *
47  * - Programatic Error Handling. libzfs_core communicates errors with
48  * defined error numbers, and doesn't print anything to stdout/stderr.
49  *
50  * - Thin Layer. libzfs_core is a thin layer, marshaling arguments
51  * to/from the kernel ioctls. There is generally a 1:1 correspondence
52  * between libzfs_core functions and ioctls to /dev/zfs.
53  *
54  * - Clear Atomicity. Because libzfs_core functions are generally 1:1
55  * with kernel ioctls, and kernel ioctls are general atomic, each
56  * libzfs_core function is atomic. For example, creating multiple
57  * snapshots with a single call to lzc_snapshot() is atomic -- it
```

new/usr/src/lib/libzfs\_core/common/libzfs\_core.c

2

```
58 * can't fail with only some of the requested snapshots created, even
59 * in the event of power loss or system crash.
60 *
61 * - Continued libzfs Support. Some higher-level operations (e.g.
62 * support for "zfs send -R") are too complicated to fit the scope of
63 * libzfs_core. This functionality will continue to live in libzfs.
64 * Where appropriate, libzfs will use the underlying atomic operations
65 * of libzfs_core. For example, libzfs may implement "zfs send -R |
66 * zfs receive" by using individual "send one snapshot", rename,
67 * destroy, and "receive one snapshot" operations in libzfs_core.
68 * /sbin/zfs and /zbin/zpool will link with both libzfs and
69 * libzfs_core. Other consumers should aim to use only libzfs_core,
70 * since that will be the supported, stable interface going forwards.
71 */
72
73 #include <libzfs_core.h>
74 #include <ctype.h>
75 #include <unistd.h>
76 #include <stdlib.h>
77 #include <string.h>
78 #include <errno.h>
79 #include <fcntl.h>
80 #include <pthread.h>
81 #include <sys/nvpair.h>
82 #include <sys/param.h>
83 #include <sys/types.h>
84 #include <sys/stat.h>
85 #include <sys/zfs_ioctl.h>
86
87 static int g_fd;
88 static pthread_mutex_t g_lock = PTHREAD_MUTEX_INITIALIZER;
89 static int g_refcount;
90
91 int
92 libzfs_core_init(void)
93 {
94     (void) pthread_mutex_lock(&g_lock);
95     if (g_refcount == 0) {
96         g_fd = open("/dev/zfs", O_RDWR);
97         if (g_fd < 0) {
98             (void) pthread_mutex_unlock(&g_lock);
99             return (errno);
100         }
101     }
102     g_refcount++;
103     (void) pthread_mutex_unlock(&g_lock);
104     return (0);
105 }
106
107 unchanged_portion_omitted
108
109 499 /*
110 500 * The simplest receive case: receive from the specified fd, creating the
111 501 * specified snapshot. Apply the specified properties a "received" properties
112 502 * (which can be overridden by locally-set properties). If the stream is a
113 503 * clone, its origin snapshot must be specified by 'origin'. The 'force'
114 504 * flag will cause the target filesystem to be rolled back or destroyed if
115 505 * necessary to receive.
116 506 *
117 507 * Return 0 on success or an errno on failure.
118 508 *
119 509 * Note: this interface does not work on dedup'd streams
120 510 * (those with DMU_BACKUP_FEATURE_DEDUP).
121 511 */
122 512 int
123 513 lzc_receive(const char *snapname, nvlist_t *props, const char *origin,
124 514             boolean_t force, int fd)
```

```

515 {
516     /*
517      * The receive ioctl is still legacy, so we need to construct our own
518      * zfs_cmd_t rather than using zpsc_ioctl().
519      */
520     zfs_cmd_t zc = { 0 };
521     char *atp;
522     char *packed = NULL;
523     size_t size;
524     dmu_replay_record_t drr;
525     int error;

527     ASSERT3S(g_refcount, >, 0);

529     /* zc_name is name of containing filesystem */
530     (void) strcpy(zc.zc_name, snapname, sizeof (zc.zc_name));
531     atp = strchr(zc.zc_name, '@');
532     if (atp == NULL)
533         return (EINVAL);
534     *atp = '\0';

536     /* if the fs does not exist, try its parent. */
537     if (!lzc_exists(zc.zc_name)) {
538         char *slashp = strrchr(zc.zc_name, '/');
539         if (slashp == NULL)
540             return (ENOENT);
541         *slashp = '\0';

543     }

545     /* zc_value is full name of the snapshot to create */
546     (void) strcpy(zc.zc_value, snapname, sizeof (zc.zc_value));

548     if (props != NULL) {
549         /* zc_nvlist_src is props to set */
550         packed = fvnlist_pack(props, &size);
551         zc.zc_nvlist_src = (uint64_t)(uintptr_t)packed;
552         zc.zc_nvlist_src_size = size;
553     }

555     /* zc_string is name of clone origin (if DRR_FLAG_CLONE) */
556     if (origin != NULL)
557         (void) strcpy(zc.zc_string, origin, sizeof (zc.zc_string));

559     /* zc_begin_record is non-byteswapped BEGIN record */
560     error = recv_read(fd, &drr, sizeof (drr));
561     if (error != 0)
562         goto out;
563     zc.zc_begin_record = drr.drr_u.drr_begin;

565     /* zc_cookie is fd to read from */
566     zc.zc_cookie = fd;

568     /* zc_guid is force flag */
569     zc.zc_guid = force;

571     /* zc_cleanup_fd is unused */
572     zc.zc_cleanup_fd = -1;

574     error = ioctl(g_fd, ZFS_IOC_RECV, &zc);
575     if (error != 0)
576         error = errno;

578 out:
579     if (packed != NULL)
580         fvnlist_pack_free(packed, size);

```

```

581         free((void*)(uintptr_t)zc.zc_nvlist_dst);
582         return (error);
583     }

585     /*
586      * Roll back this filesystem or volume to its most recent snapshot.
587      * If snapnamebuf is not NULL, it will be filled in with the name
588      * of the most recent snapshot.
589      *
590      * Return 0 on success or an errno on failure.
591      */
592     int
593     lzc_rollback(const char *fsname, char *snapnamebuf, int snapnamelen)
594     {
595         nvlist_t *args;
596         nvlist_t *result;
597         int err;

599         args = fvnlist_alloc();
600         err = lzc_ioctl(ZFS_IOC_ROLLBACK, fsname, args, &result);
601         nvlist_free(args);
602         if (err == 0 && snapnamebuf != NULL) {
603             const char *snapname = fvnlist_lookup_string(result, "target");
604             (void) strcpy(snapnamebuf, snapname, snapnamelen);
605         }
606         return (err);
607     }

```

\_\_\_\_\_unchanged\_portion\_omitted\_\_\_\_\_

new/usr/src/lib/libzfs\_core/common/libzfs\_core.h

1

```
*****
2130 Mon Aug 12 01:16:05 2013
new/usr/src/lib/libzfs_core/common/libzfs_core.h
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */

22 /*
23  * Copyright (c) 2013 by Delphix. All rights reserved.
24  * Copyright (c) 2012 by Delphix. All rights reserved.
25 */

26 #ifndef _LIBZFS_CORE_H
27 #define _LIBZFS_CORE_H

28 #include <libnvpair.h>
29 #include <sys/param.h>
30 #include <sys/types.h>
31 #include <sys/fs/zfs.h>

32 #ifdef __cplusplus
33 extern "C" {
34 #endif

35 int libzfs_core_init(void);
36 void libzfs_core_fini(void);

37
38 int lzc_snapshot(nvlist_t *snaps, nvlist_t *props, nvlist_t **errlist);
39 int lzc_create(const char *fsname, dmu_objset_type_t type, nvlist_t *props);
40 int lzc_clone(const char *fsname, const char *origin, nvlist_t *props);
41 int lzc_destroy_snaps(nvlist_t *snaps, boolean_t defer, nvlist_t **errlist);

42 int lzc_snaprange_space(const char *firstsnap, const char *lastsnap,
43 uint64_t *usedp);

44 int lzc_hold(nvlist_t *holds, int cleanup_fd, nvlist_t **errlist);
45 int lzc_release(nvlist_t *holds, nvlist_t **errlist);
46 int lzc_get_holds(const char *snapname, nvlist_t **holdsp);

47
48 int lzc_send(const char *snapname, const char *fromsnap, int fd);
49 int lzc_receive(const char *snapname, nvlist_t *props, const char *origin,
50 boolean_t force, int fd);
51 int lzc_send_space(const char *snapname, const char *fromsnap,
52 uint64_t *result);
```

new/usr/src/lib/libzfs\_core/common/libzfs\_core.h

2

```
59 boolean_t lzc_exists(const char *dataset);

60 int lzc_rollback(const char *fsname, char *snapnamebuf, int snapnamelen);

61
62 #ifdef __cplusplus
63 }
64 }

unchanged_portion_omitted
```

new/usr/src/lib/libzfs\_core/common/mapfile-vers

1

\*\*\*\*\*

1593 Mon Aug 12 01:16:06 2013

new/usr/src/lib/libzfs\_core/common/mapfile-vers

3996 want a libzfs\_core API to rollback to latest snapshot

Reviewed by: Christopher Siden <christopher.siden@delphix.com>

Reviewed by: Adam Leventhal <ahl@delphix.com>

Reviewed by: George Wilson <george.wilson@delphix.com>

\*\*\*\*\*

```
1 #
2 # CDDL HEADER START
3 #
4 # The contents of this file are subject to the terms of the
5 # Common Development and Distribution License (the "License").
6 # You may not use this file except in compliance with the License.
7 #
8 # You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9 # or http://www.opensolaris.org/os/licensing.
10 # See the License for the specific language governing permissions
11 # and limitations under the License.
12 #
13 # When distributing Covered Code, include this CDDL HEADER in each
14 # file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 # If applicable, add the following below this CDDL HEADER, with the
16 # fields enclosed by brackets "[]" replaced with your own identifying
17 # information: Portions Copyright [yyyy] [name of copyright owner]
18 #
19 # CDDL HEADER END
20 #
21 # Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
22 # Copyright (c) 2013 by Delphix. All rights reserved.
23 # Copyright (c) 2012 by Delphix. All rights reserved.
24 #
25 # MAPFILE HEADER START
26 #
27 # WARNING: STOP NOW. DO NOT MODIFY THIS FILE.
28 # Object versioning must comply with the rules detailed in
29 #     usr/src/lib/README.mapfiles
30 #
31 # You should not be making modifications here until you've read the most current
32 # copy of that file. If you need help, contact a gatekeeper for guidance.
33 #
34 # MAPFILE HEADER END
35 #
36
37 $mapfile_version 2
38
39 SYMBOL_VERSION ILLUMOS_0.1 {
40     global:
41
42         libzfs_core_fini;
43         libzfs_core_init;
44         lzc_clone;
45         lzc_create;
46         lzc_destroy_snaps;
47         lzc_exists;
48         lzc_get_holds;
49         lzc_hold;
50         lzc_receive;
51         lzc_release;
52         lzc_rollback;
53         lzc_send;
54         lzc_send_space;
55         lzc_snaprange_space;
56         lzc_snapshot;
```

new/usr/src/lib/libzfs\_core/common/mapfile-vers

2

```
58     local:
59         *;
60 };
_____unchanged_portion_omitted_
```

```

*****
83061 Mon Aug 12 01:16:07 2013
new/usr/src/uts/common/fs/zfs/dsl_dataset.c
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
_____unchanged_portion_omitted_____

1722 typedef struct dsl_dataset_rollback_arg {
1723     const char *ddra_fsname;
1724     void *ddra_owner;
1725     nvlist_t *ddra_result;
1726 } dsl_dataset_rollback_arg_t;
_____unchanged_portion_omitted_____

1790 static void
1791 dsl_dataset_rollback_sync(void *arg, dmu_tx_t *tx)
1792 {
1793     dsl_dataset_rollback_arg_t *ddra = arg;
1794     dsl_pool_t *dp = dmu_tx_pool(tx);
1795     dsl_dataset_t *ds, *clone;
1796     uint64_t cloneobj;
1797     char namebuf[ZFS_MAXNAMELEN];

1799     VERIFY0(dsl_dataset_hold(dp, ddra->ddra_fsname, FTAG, &ds));

1801     dsl_dataset_name(ds->ds_prev, namebuf);
1802     fnvlist_add_string(ddra->ddra_result, "target", namebuf);

1804     cloneobj = dsl_dataset_create_sync(ds->ds_dir, "%rollback",
1805     ds->ds_prev, DS_CREATE_FLAG_NODIRTY, kcred, tx);

1807     VERIFY0(dsl_dataset_hold_obj(dp, cloneobj, FTAG, &clone));

1809     dsl_dataset_clone_swap_sync_impl(clone, ds, tx);
1810     dsl_dataset_zero_zil(ds, tx);

1812     dsl_destroy_head_sync_impl(clone, tx);

1814     dsl_dataset_rele(clone, FTAG);
1815     dsl_dataset_rele(ds, FTAG);
1816 }

1818 /*
1819  * Rolls back the given filesystem or volume to the most recent snapshot.
1820  * The name of the most recent snapshot will be returned under key "target"
1821  * in the result nvlist.
1822  *
1823  * If owner != NULL:
1824  * - The existing dataset MUST be owned by the specified owner at entry
1825  * - Upon return, dataset will still be held by the same owner, whether we
1826  *   succeed or not.
1827  *
1828  * This mode is required any time the existing filesystem is mounted. See
1829  * notes above zfs_suspend_fs() for further details.
1830  */
1831 int
1832 dsl_dataset_rollback(const char *fsname, void *owner, nvlist_t *result)
1833 dsl_dataset_rollback(const char *fsname, void *owner)
1834 {
1834     dsl_dataset_rollback_arg_t ddra;

1836     ddra.ddra_fsname = fsname;

```

```

1837     ddra.ddra_owner = owner;
1838     ddra.ddra_result = result;

1840     return (dsl_sync_task(fsname, dsl_dataset_rollback_check,
1841     dsl_dataset_rollback_sync, &ddra, 1));
1832     dsl_dataset_rollback_sync, (void *)&ddra, 1));
1842 }
_____unchanged_portion_omitted_____

```

```

*****
10302 Mon Aug 12 01:16:08 2013
new/usr/src/uts/common/fs/zfs/sys/dsl_dataset.h
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
_____unchanged_portion_omitted_____

167 /*
168 * The max length of a temporary tag prefix is the number of hex digits
169 * required to express UINT64_MAX plus one for the hyphen.
170 */
171 #define MAX_TAG_PREFIX_LEN      17

173 #define dsl_dataset_is_snapshot(ds) \
174     ((ds)->ds_phys->ds_num_children != 0)

176 #define DS_UNIQUE_IS_ACCURATE(ds) \
177     (((ds)->ds_phys->ds_flags & DS_FLAG_UNIQUE_ACCURATE) != 0)

179 int dsl_dataset_hold(struct dsl_pool *dp, const char *name, void *tag,
180     dsl_dataset_t **dsp);
181 int dsl_dataset_hold_obj(struct dsl_pool *dp, uint64_t dsobj, void *tag,
182     dsl_dataset_t **);
183 void dsl_dataset_rele(dsl_dataset_t *ds, void *tag);
184 int dsl_dataset_own(struct dsl_pool *dp, const char *name,
185     void *tag, dsl_dataset_t **dsp);
186 int dsl_dataset_own_obj(struct dsl_pool *dp, uint64_t dsobj,
187     void *tag, dsl_dataset_t **dsp);
188 void dsl_dataset_disown(dsl_dataset_t *ds, void *tag);
189 void dsl_dataset_name(dsl_dataset_t *ds, char *name);
190 boolean_t dsl_dataset_tryown(dsl_dataset_t *ds, void *tag);
191 uint64_t dsl_dataset_create_sync(dsl_dir_t *pds, const char *lastname,
192     dsl_dataset_t *origin, uint64_t flags, cred_t *, dmu_tx_t *);
193 uint64_t dsl_dataset_create_sync_dd(dsl_dir_t *dd, dsl_dataset_t *origin,
194     uint64_t flags, dmu_tx_t *tx);
195 int dsl_dataset_snapshot(nvlist_t *snaps, nvlist_t *props, nvlist_t *errors);
196 int dsl_dataset_promote(const char *name, char *conflsnap);
197 int dsl_dataset_clone_swap(dsl_dataset_t *clone, dsl_dataset_t *origin_head,
198     boolean_t force);
199 int dsl_dataset_rename_snapshot(const char *fsname,
200     const char *oldsnapname, const char *newsnapname, boolean_t recursive);
201 int dsl_dataset_snapshot_tmp(const char *fsname, const char *snapname,
202     minor_t cleanup_minor, const char *htag);

204 blkptr_t *dsl_dataset_get_blkptr(dsl_dataset_t *ds);
205 void dsl_dataset_set_blkptr(dsl_dataset_t *ds, blkptr_t *bp, dmu_tx_t *tx);

207 spa_t *dsl_dataset_get_spa(dsl_dataset_t *ds);

209 boolean_t dsl_dataset_modified_since_snap(dsl_dataset_t *ds,
210     dsl_dataset_t *snap);

212 void dsl_dataset_sync(dsl_dataset_t *os, zio_t *zio, dmu_tx_t *tx);

214 void dsl_dataset_block_born(dsl_dataset_t *ds, const blkptr_t *bp,
215     dmu_tx_t *tx);
216 int dsl_dataset_block_kill(dsl_dataset_t *ds, const blkptr_t *bp,
217     dmu_tx_t *tx, boolean_t async);
218 boolean_t dsl_dataset_block_freeable(dsl_dataset_t *ds, const blkptr_t *bp,
219     uint64_t blk_birth);
220 uint64_t dsl_dataset_prev_snap_txg(dsl_dataset_t *ds);

222 void dsl_dataset_dirty(dsl_dataset_t *ds, dmu_tx_t *tx);

```

```

223 void dsl_dataset_stats(dsl_dataset_t *os, nvlist_t *nv);
224 void dsl_dataset_fast_stat(dsl_dataset_t *ds, dmu_objset_stats_t *stat);
225 void dsl_dataset_space(dsl_dataset_t *ds,
226     uint64_t *refdbbytesp, uint64_t *availbytesp,
227     uint64_t *usedobjsp, uint64_t *availobjsp);
228 uint64_t dsl_dataset_fsid_guid(dsl_dataset_t *ds);
229 int dsl_dataset_space_written(dsl_dataset_t *oldsnap, dsl_dataset_t *new,
230     uint64_t *usedp, uint64_t *compp, uint64_t *uncompp);
231 int dsl_dataset_space_wouldfree(dsl_dataset_t *firstsnap, dsl_dataset_t *last,
232     uint64_t *usedp, uint64_t *compp, uint64_t *uncompp);
233 boolean_t dsl_dataset_is_dirty(dsl_dataset_t *ds);

235 int dsl_dsobj_to_dsname(char *pname, uint64_t obj, char *buf);

237 int dsl_dataset_check_quota(dsl_dataset_t *ds, boolean_t check_quota,
238     uint64_t asize, uint64_t inflight, uint64_t *used,
239     uint64_t *ref_rsrv);
240 int dsl_dataset_set_refquota(const char *dsname, zprop_source_t source,
241     uint64_t quota);
242 int dsl_dataset_set_refreservation(const char *dsname, zprop_source_t source,
243     uint64_t reservation);

245 boolean_t dsl_dataset_is_before(dsl_dataset_t *later, dsl_dataset_t *earlier);
246 void dsl_dataset_long_hold(dsl_dataset_t *ds, void *tag);
247 void dsl_dataset_long_rele(dsl_dataset_t *ds, void *tag);
248 boolean_t dsl_dataset_long_held(dsl_dataset_t *ds);

250 int dsl_dataset_clone_swap_check_impl(dsl_dataset_t *clone,
251     dsl_dataset_t *origin_head, boolean_t force, void *owner, dmu_tx_t *tx);
252 void dsl_dataset_clone_swap_sync_impl(dsl_dataset_t *clone,
253     dsl_dataset_t *origin_head, dmu_tx_t *tx);
254 int dsl_dataset_snapshot_check_impl(dsl_dataset_t *ds, const char *snapname,
255     dmu_tx_t *tx, boolean_t recv);
256 void dsl_dataset_snapshot_sync_impl(dsl_dataset_t *ds, const char *snapname,
257     dmu_tx_t *tx);

259 void dsl_dataset_remove_from_next_clones(dsl_dataset_t *ds, uint64_t obj,
260     dmu_tx_t *tx);
261 void dsl_dataset_recalc_head_uniq(dsl_dataset_t *ds);
262 int dsl_dataset_get_snapname(dsl_dataset_t *ds);
263 int dsl_dataset_snap_lookup(dsl_dataset_t *ds, const char *name,
264     uint64_t *value);
265 int dsl_dataset_snap_remove(dsl_dataset_t *ds, const char *name, dmu_tx_t *tx);
266 void dsl_dataset_set_refreservation_sync_impl(dsl_dataset_t *ds,
267     zprop_source_t source, uint64_t value, dmu_tx_t *tx);
268 int dsl_dataset_rollback(const char *fsname, void *owner, nvlist_t *result);
268 int dsl_dataset_rollback(const char *fsname, void *owner);

270 #ifdef ZFS_DEBUG
271 #define dprintf_ds(ds, fmt, ...) do { \
272     if (zfs_flags & ZFS_DEBUG_DPRINTF) { \
273         char * _ds_name = kmem_alloc(MAXNAMELEN, KM_SLEEP); \
274         dsl_dataset_name(ds, _ds_name); \
275         dprintf("ds=%s " fmt, _ds_name, __VA_ARGS__); \
276         kmem_free(_ds_name, MAXNAMELEN); \
277     } \
278     _NOTE(CONSTCOND) } while (0)
279 #else
280 #define dprintf_ds(dd, fmt, ...)
281 #endif

283 #ifdef __cplusplus
284 }
285 #endif

287 #endif /* _SYS_DSL_DATASET_H */

```

new/usr/src/uts/common/fs/zfs/zfs\_ioctl.c

1

```
*****
144478 Mon Aug 12 01:16:10 2013
new/usr/src/uts/common/fs/zfs/zfs_ioctl.c
3996 want a libzfs_core API to rollback to latest snapshot
Reviewed by: Christopher Siden <christopher.siden@delphix.com>
Reviewed by: Adam Leventhal <ahl@delphix.com>
Reviewed by: George Wilson <george.wilson@delphix.com>
*****
_____unchanged_portion_omitted_____

3495 /*
3496 * fsname is name of dataset to rollback (to most recent snapshot)
3497 * inputs:
3497 *   zc_name      name of dataset to rollback (to most recent snapshot)
3497 *
3498 * innvl is not used.
3499 *
3500 * outnvl: "target" -> name of most recent snapshot
3501 * }
3499 * outputs:      none
3502 */
3503 /* ARGSUSED */
3504 static int
3505 zfs_ioc_rollback(const char *fsname, nvlist_t *args, nvlist_t *outnvl)
3506 zfs_ioc_rollback(zfs_cmd_t *zc)
3506 {
3507     zfsvfs_t *zfsvfs;
3508     int error;

3510     if (getzfsvfs(fsname, &zfsvfs) == 0) {
3507         if (getzfsvfs(zc->zc_name, &zfsvfs) == 0) {
3511             error = zfs_suspend_fs(zfsvfs);
3512             if (error == 0) {
3513                 int resume_err;

3515                 error = dsl_dataset_rollback(fsname, zfsvfs, outnvl);
3516                 resume_err = zfs_resume_fs(zfsvfs, fsname);
3517                 error = dsl_dataset_rollback(zc->zc_name, zfsvfs);
3518                 resume_err = zfs_resume_fs(zfsvfs, zc->zc_name);
3519                 error = error ? error : resume_err;
3517             }
3518             VFS_RELE(zfsvfs->z_vfs);
3520         } else {
3521             error = dsl_dataset_rollback(fsname, NULL, outnvl);
3518             error = dsl_dataset_rollback(zc->zc_name, NULL);
3522         }
3523         return (error);
3524     }

_____unchanged_portion_omitted_____

5287 static void
5288 zfs_ioctl_init(void)
5289 {
5290     zfs_ioctl_register("snapshot", ZFS_IOC_SNAPSHOT,
5291         zfs_ioc_snapshot, zfs_secpolicy_snapshot, POOL_NAME,
5292         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);

5294     zfs_ioctl_register("log_history", ZFS_IOC_LOG_HISTORY,
5295         zfs_ioc_log_history, zfs_secpolicy_log_history, NO_NAME,
5296         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_FALSE, B_FALSE);

5298     zfs_ioctl_register("space_snaps", ZFS_IOC_SPACE_SNAPS,
5299         zfs_ioc_space_snaps, zfs_secpolicy_read, DATASET_NAME,
5300         POOL_CHECK_SUSPENDED, B_FALSE, B_FALSE);

5302     zfs_ioctl_register("send", ZFS_IOC_SEND_NEW,
```

new/usr/src/uts/common/fs/zfs/zfs\_ioctl.c

2

```
5303         zfs_ioc_send_new, zfs_secpolicy_send_new, DATASET_NAME,
5304         POOL_CHECK_SUSPENDED, B_FALSE, B_FALSE);

5306     zfs_ioctl_register("send_space", ZFS_IOC_SEND_SPACE,
5307         zfs_ioc_send_space, zfs_secpolicy_read, DATASET_NAME,
5308         POOL_CHECK_SUSPENDED, B_FALSE, B_FALSE);

5310     zfs_ioctl_register("create", ZFS_IOC_CREATE,
5311         zfs_ioc_create, zfs_secpolicy_create_clone, DATASET_NAME,
5312         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);

5314     zfs_ioctl_register("clone", ZFS_IOC_CLONE,
5315         zfs_ioc_clone, zfs_secpolicy_create_clone, DATASET_NAME,
5316         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);

5318     zfs_ioctl_register("destroy_snaps", ZFS_IOC_DESTROY_SNAPS,
5319         zfs_ioc_destroy_snaps, zfs_secpolicy_destroy_snaps, POOL_NAME,
5320         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);

5322     zfs_ioctl_register("hold", ZFS_IOC_HOLD,
5323         zfs_ioc_hold, zfs_secpolicy_hold, POOL_NAME,
5324         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);
5325     zfs_ioctl_register("release", ZFS_IOC_RELEASE,
5326         zfs_ioc_release, zfs_secpolicy_release, POOL_NAME,
5327         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_TRUE, B_TRUE);

5329     zfs_ioctl_register("get_holds", ZFS_IOC_GET_HOLDS,
5330         zfs_ioc_get_holds, zfs_secpolicy_read, DATASET_NAME,
5331         POOL_CHECK_SUSPENDED, B_FALSE, B_FALSE);

5333     zfs_ioctl_register("rollback", ZFS_IOC_ROLLBACK,
5334         zfs_ioc_rollback, zfs_secpolicy_rollback, DATASET_NAME,
5335         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY, B_FALSE, B_TRUE);

5337     /* IOCTLS that use the legacy function signature */

5339     zfs_ioctl_register_legacy(ZFS_IOC_POOL_FREEZE, zfs_ioc_pool_freeze,
5340         zfs_secpolicy_config, NO_NAME, B_FALSE, POOL_CHECK_READONLY);

5342     zfs_ioctl_register_pool(ZFS_IOC_POOL_CREATE, zfs_ioc_pool_create,
5343         zfs_secpolicy_config, B_TRUE, POOL_CHECK_NONE);
5344     zfs_ioctl_register_pool_modify(ZFS_IOC_POOL_SCAN,
5345         zfs_ioc_pool_scan);
5346     zfs_ioctl_register_pool_modify(ZFS_IOC_POOL_UPGRADE,
5347         zfs_ioc_pool_upgrade);
5348     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_ADD,
5349         zfs_ioc_vdev_add);
5350     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_REMOVE,
5351         zfs_ioc_vdev_remove);
5352     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_SET_STATE,
5353         zfs_ioc_vdev_set_state);
5354     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_ATTACH,
5355         zfs_ioc_vdev_attach);
5356     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_DETACH,
5357         zfs_ioc_vdev_detach);
5358     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_SETPATH,
5359         zfs_ioc_vdev_setpath);
5360     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_SETFRU,
5361         zfs_ioc_vdev_setfru);
5362     zfs_ioctl_register_pool_modify(ZFS_IOC_POOL_SET_PROPS,
5363         zfs_ioc_pool_set_props);
5364     zfs_ioctl_register_pool_modify(ZFS_IOC_VDEV_SPLIT,
5365         zfs_ioc_vdev_split);
5366     zfs_ioctl_register_pool_modify(ZFS_IOC_POOL_REGUID,
5367         zfs_ioc_pool_reguid);
```



```

5369     zfs_ioctl_register_pool_meta(ZFS_IOC_POOL_CONFIGS,
5370         zfs_ioc_pool_configs, zfs_secpolicy_none);
5371     zfs_ioctl_register_pool_meta(ZFS_IOC_POOL_TRYIMPORT,
5372         zfs_ioc_pool_tryimport, zfs_secpolicy_config);
5373     zfs_ioctl_register_pool_meta(ZFS_IOC_INJECT_FAULT,
5374         zfs_ioc_inject_fault, zfs_secpolicy_inject);
5375     zfs_ioctl_register_pool_meta(ZFS_IOC_CLEAR_FAULT,
5376         zfs_ioc_clear_fault, zfs_secpolicy_inject);
5377     zfs_ioctl_register_pool_meta(ZFS_IOC_INJECT_LIST_NEXT,
5378         zfs_ioc_inject_list_next, zfs_secpolicy_inject);

5380     /*
5381      * pool destroy, and export don't log the history as part of
5382      * zfsdev_ioctl, but rather zfs_ioc_pool_export
5383      * does the logging of those commands.
5384      */
5385     zfs_ioctl_register_pool(ZFS_IOC_POOL_DESTROY, zfs_ioc_pool_destroy,
5386         zfs_secpolicy_config, B_FALSE, POOL_CHECK_NONE);
5387     zfs_ioctl_register_pool(ZFS_IOC_POOL_EXPORT, zfs_ioc_pool_export,
5388         zfs_secpolicy_config, B_FALSE, POOL_CHECK_NONE);

5390     zfs_ioctl_register_pool(ZFS_IOC_POOL_STATS, zfs_ioc_pool_stats,
5391         zfs_secpolicy_read, B_FALSE, POOL_CHECK_NONE);
5392     zfs_ioctl_register_pool(ZFS_IOC_POOL_GET_PROPS, zfs_ioc_pool_get_props,
5393         zfs_secpolicy_read, B_FALSE, POOL_CHECK_NONE);

5395     zfs_ioctl_register_pool(ZFS_IOC_ERROR_LOG, zfs_ioc_error_log,
5396         zfs_secpolicy_inject, B_FALSE, POOL_CHECK_SUSPENDED);
5397     zfs_ioctl_register_pool(ZFS_IOC_DSOBJ_TO_DSNAME,
5398         zfs_ioc_dsobj_to_dsname,
5399         zfs_secpolicy_diff, B_FALSE, POOL_CHECK_SUSPENDED);
5400     zfs_ioctl_register_pool(ZFS_IOC_POOL_GET_HISTORY,
5401         zfs_ioc_pool_get_history,
5402         zfs_secpolicy_config, B_FALSE, POOL_CHECK_SUSPENDED);

5404     zfs_ioctl_register_pool(ZFS_IOC_POOL_IMPORT, zfs_ioc_pool_import,
5405         zfs_secpolicy_config, B_TRUE, POOL_CHECK_NONE);

5407     zfs_ioctl_register_pool(ZFS_IOC_CLEAR, zfs_ioc_clear,
5408         zfs_secpolicy_config, B_TRUE, POOL_CHECK_SUSPENDED);
5409     zfs_ioctl_register_pool(ZFS_IOC_POOL_REOPEN, zfs_ioc_pool_reopen,
5410         zfs_secpolicy_config, B_TRUE, POOL_CHECK_SUSPENDED);

5412     zfs_ioctl_register_dataset_read(ZFS_IOC_SPACE_WRITTEN,
5413         zfs_ioc_space_written);
5414     zfs_ioctl_register_dataset_read(ZFS_IOC_OBJSET_RECVD_PROPS,
5415         zfs_ioc_objset_recvd_props);
5416     zfs_ioctl_register_dataset_read(ZFS_IOC_NEXT_OBJ,
5417         zfs_ioc_next_obj);
5418     zfs_ioctl_register_dataset_read(ZFS_IOC_GET_FSACL,
5419         zfs_ioc_get_fsacl);
5420     zfs_ioctl_register_dataset_read(ZFS_IOC_OBJSET_STATS,
5421         zfs_ioc_objset_stats);
5422     zfs_ioctl_register_dataset_read(ZFS_IOC_OBJSET_ZPLPROPS,
5423         zfs_ioc_objset_zplprops);
5424     zfs_ioctl_register_dataset_read(ZFS_IOC_DATASET_LIST_NEXT,
5425         zfs_ioc_dataset_list_next);
5426     zfs_ioctl_register_dataset_read(ZFS_IOC_SNAPSHOT_LIST_NEXT,
5427         zfs_ioc_snapshot_list_next);
5428     zfs_ioctl_register_dataset_read(ZFS_IOC_SEND_PROGRESS,
5429         zfs_ioc_send_progress);

5431     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_DIFF,
5432         zfs_ioc_diff, zfs_secpolicy_diff);
5433     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_OBJ_TO_STATS,
5434         zfs_ioc_obj_to_stats, zfs_secpolicy_diff);

```

```

5435     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_OBJ_TO_PATH,
5436         zfs_ioc_obj_to_path, zfs_secpolicy_diff);
5437     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_USERSPACE_ONE,
5438         zfs_ioc_userspace_one, zfs_secpolicy_userspace_one);
5439     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_USERSPACE_MANY,
5440         zfs_ioc_userspace_many, zfs_secpolicy_userspace_many);
5441     zfs_ioctl_register_dataset_read_secpolicy(ZFS_IOC_SEND,
5442         zfs_ioc_send, zfs_secpolicy_send);

5444     zfs_ioctl_register_dataset_modify(ZFS_IOC_SET_PROP, zfs_ioc_set_prop,
5445         zfs_secpolicy_none);
5446     zfs_ioctl_register_dataset_modify(ZFS_IOC_DESTROY, zfs_ioc_destroy,
5447         zfs_secpolicy_destroy);
5448     zfs_ioctl_register_dataset_modify(ZFS_IOC_ROLLBACK, zfs_ioc_rollback,
5449         zfs_secpolicy_rollback);
5450     zfs_ioctl_register_dataset_modify(ZFS_IOC_RENAME, zfs_ioc_rename,
5451         zfs_secpolicy_rename);
5452     zfs_ioctl_register_dataset_modify(ZFS_IOC_RECV, zfs_ioc_recv,
5453         zfs_secpolicy_recv);
5454     zfs_ioctl_register_dataset_modify(ZFS_IOC_PROMOTE, zfs_ioc_promote,
5455         zfs_secpolicy_promote);
5456     zfs_ioctl_register_dataset_modify(ZFS_IOC_INHERIT_PROP,
5457         zfs_ioc_inherit_prop, zfs_secpolicy_inherit_prop);
5458     zfs_ioctl_register_dataset_modify(ZFS_IOC_SET_FSACL, zfs_ioc_set_fsacl,
5459         zfs_secpolicy_set_fsacl);

5459     zfs_ioctl_register_dataset_nolog(ZFS_IOC_SHARE, zfs_ioc_share,
5460         zfs_secpolicy_share, POOL_CHECK_NONE);
5461     zfs_ioctl_register_dataset_nolog(ZFS_IOC_SMB_ACL, zfs_ioc_smb_acl,
5462         zfs_secpolicy_smb_acl, POOL_CHECK_NONE);
5463     zfs_ioctl_register_dataset_nolog(ZFS_IOC_USERSPACE_UPGRADE,
5464         zfs_ioc_userspace_upgrade, zfs_secpolicy_userspace_upgrade,
5465         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY);
5466     zfs_ioctl_register_dataset_nolog(ZFS_IOC_TMP_SNAPSHOT,
5467         zfs_ioc_tmp_snapshot, zfs_secpolicy_tmp_snapshot,
5468         POOL_CHECK_SUSPENDED | POOL_CHECK_READONLY);
5469 }

```

unchanged portion omitted