

```
new/usr/src/uts/common/fs/dev/sdev_profile.c
```

```
*****
26638 Mon Dec 1 22:44:39 2014
new/usr/src/uts/common/fs/dev/sdev_profile.c
5360 Race condition in devfs upgrades reader to writer incidentally and causes p
*****
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 2014 Coraid, Inc.
24 * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
25 */
27 /*
28 * This file implements /dev filesystem operations for non-global
29 * instances. Three major entry points:
30 * devname_profile_update()
31 *   Update matching rules determining which names to export
32 * prof_readdir()
33 *   Return the list of exported names
34 * prof_lookup()
35 *   Implements lookup
36 */
38 #include <sys/types.h>
39 #include <sys/param.h>
40 #include <sys/sysmacros.h>
41 #include <sys/vnode.h>
42 #include <sys/uio.h>
43 #include <sys/dirent.h>
44 #include <sys/pathname.h>
45 #include <sys/fs/dv_node.h>
46 #include <sys/fs/sdev_impl.h>
47 #include <sys/sunndi.h>
48 #include <sys/modctl.h>
50 enum {
51     PROFILE_TYPE_INCLUDE,
52     PROFILE_TYPE_EXCLUDE,
53     PROFILE_TYPE_MAP,
54     PROFILE_TYPE_SYMLINK
55 };
56 unchanged_portion_omitted_
658 /*
659 * Return True if directory cache is out of date and should be updated.
660 */
661 static boolean_t
```

```
1
```

```
new/usr/src/uts/common/fs/dev/sdev_profile.c
```

```
662 prof_dev_needupdate(sdev_node_t *ddv)
663 {
664     sdev_node_t *gdir = ddv->sdev_origin;
665
666     /*
667      * Caller can have either reader or writer lock
668      */
669     ASSERT(RW_LOCK_HELD(&ddv->sdev_contents));
670
671     /*
672      * We need to rebuild the directory content if
673      * - ddv is not in a SDEV_ZOMBIE state
674      * - SDEV_BUILD is set OR
675      * - The device tree generation number has changed OR
676      * - The corresponding /dev namespace has been updated
677      */
678     return ((ddv->sdev_state != SDEV_ZOMBIE) &&
679            (((ddv->sdev_flags & SDEV_BUILD) != 0) ||
680             (ddv->sdev_devtree_gen != devtree_gen)) ||
681            ((gdir != NULL) &&
682             (ddv->sdev_ldir_gen != gdir->sdev_gdir_gen)));
683 }
685 /*
686  * Build directory vnodes based on the profile and the global
687  * dev instance.
688 */
689 void
690 prof_filldir(sdev_node_t *ddv)
691 {
692     sdev_node_t *gdir;
693     int firsttime = 1;
694     struct sdev_node *gdir = ddv->sdev_origin;
695
696     ASSERT(RW_READ_HELD(&ddv->sdev_contents));
697     if (!prof_dev_needupdate(ddv)) {
698         ASSERT(RW_READ_HELD(&ddv->sdev_contents));
699         return;
700     }
701     /*
702      * Upgrade to writer lock
703      * We need to rebuild the directory content if
704      * - SDEV_BUILD is set
705      * - The device tree generation number has changed
706      * - The corresponding /dev namespace has been updated
707      */
708     if (rw_tryupgrade(&ddv->sdev_contents) == 0) {
709         /*
710          * We need to drop the read lock and re-acquire it as a
711          * write lock. While we do this the condition may change so we
712          * need to re-check condition.
713          * NOTE: if device becomes a zombie, prof_dev_needupdate returns
714          * false, so nothing we will just return in this case.
715        */
716        check_build:
717        if ((ddv->sdev_flags & SDEV_BUILD) == 0 &&
718            ddv->sdev_devtree_gen == devtree_gen &&
719            (gdir == NULL || ddv->sdev_ldir_gen
720             == gdir->sdev_gdir_gen))
721            return; /* already up to date */
722
723        /*
724         * We may have become a zombie (across a try)
725         */
726        if (ddv->sdev_state == SDEV_ZOMBIE)
727            return;
```

```
2
```

```
686     if (firsttime && rw_tryupgrade(&ddv->sdev_contents) == 0) {
711         rw_exit(&ddv->sdev_contents);
712         firsttime = 0;
713         rw_enter(&ddv->sdev_contents, RW_WRITER);
714         if (!prof_dev_needupdate(ddv)) {
715             /* Downgrade back to the read lock before returning */
716             rw_downgrade(&ddv->sdev_contents);
717             return;
718         }
719         /* At this point we should have a write lock */
720         ASSERT(RW_WRITE_HELD(&ddv->sdev_contents));
722         sdcmn_err10(("devtree_gen (%s): %ld -> %ld\n",
723                     ddv->sdev_path, ddv->sdev_devtree_gen, devtree_gen));
725         gdir = ddv->sdev_origin;
727         if (gdir != NULL)
728             if (gdir)
729                 sdcmn_err10(("sdev_dir_gen (%s): %ld -> %ld\n",
730                             ddv->sdev_path, ddv->sdev_ldir_gen,
731                             gdir->sdev_gdir_gen));
732         /* update flags and generation number so next filldir is quick */
733         if ((ddv->sdev_flags & SDEV_BUILD) == SDEV_BUILD) {
734             ddv->sdev_flags &= ~SDEV_BUILD;
735         }
736         ddv->sdev_devtree_gen = devtree_gen;
737         if (gdir != NULL)
738             if (gdir)
739                 ddv->sdev_ldir_gen = gdir->sdev_gdir_gen;
740         prof_make_symlinks(ddv);
741         prof_make_maps(ddv);
742         prof_make_names(ddv);
743         rw_downgrade(&ddv->sdev_contents);
744 }
```

unchanged portion omitted