

```

*****
110792 Mon Jan 14 13:16:35 2019
new/usr/src/uts/common/fs/vnode.c
10084 fop_open() doesn't need to check for a NULL vnode
*****
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) 1988, 2010, Oracle and/or its affiliates. All rights reserved.
24  * Copyright (c) 2018, Joyent, Inc.
25  * Copyright 2017, Joyent, Inc.
26  * Copyright 2016 Nexenta Systems, Inc. All rights reserved.
27  * Copyright (c) 2011, 2017 by Delphix. All rights reserved.
28 */

29 /*      Copyright (c) 1983, 1984, 1985, 1986, 1987, 1988, 1989 AT&T      */
30 /*      All Rights Reserved      */

32 /*
33  * University Copyright- Copyright (c) 1982, 1986, 1988
34  * The Regents of the University of California
35  * All Rights Reserved
36  *
37  * University Acknowledgment- Portions of this document are derived from
38  * software developed by the University of California, Berkeley, and its
39  * contributors.
40 */

42 #include <sys/types.h>
43 #include <sys/param.h>
44 #include <sys/t_lock.h>
45 #include <sys/errno.h>
46 #include <sys/cred.h>
47 #include <sys/user.h>
48 #include <sys/uid.h>
49 #include <sys/file.h>
50 #include <sys/pathname.h>
51 #include <sys/vfs.h>
52 #include <sys/vfs_opreg.h>
53 #include <sys/vnode.h>
54 #include <sys/rwstlock.h>
55 #include <sys/fem.h>
56 #include <sys/stat.h>
57 #include <sys/mode.h>
58 #include <sys/conf.h>
59 #include <sys/sysmacros.h>
60 #include <sys/cmn_err.h>

```

```

61 #include <sys/system.h>
62 #include <sys/kmem.h>
63 #include <sys/debug.h>
64 #include <c2/audit.h>
65 #include <sys/acl.h>
66 #include <sys/nbmlock.h>
67 #include <sys/fcntl.h>
68 #include <fs/fs_subr.h>
69 #include <sys/taskq.h>
70 #include <fs/fs_reparse.h>
71 #include <sys/time.h>
72 #include <sys/sdt.h>

74 /* Determine if this vnode is a file that is read-only */
75 #define ISROFILE(vp) \
76     ((vp)->v_type != VCHR && (vp)->v_type != VBLK && \
77      (vp)->v_type != VFIFO && vn_is_readonly(vp))

79 /* Tunable via /etc/system; used only by admin/install */
80 int nfs_global_client_only;

82 /*
83  * Array of vopstats_t for per-FS-type vopstats. This array has the same
84  * number of entries as and parallel to the vfssw table. (Arguably, it could
85  * be part of the vfssw table.) Once it's initialized, it's accessed using
86  * the same fstype index that is used to index into the vfssw table.
87  */
88 vopstats_t **vopstats_fstype;

90 /* vopstats initialization template used for fast initialization via bcopy() */
91 static vopstats_t *vs_templatep;

93 /* Kmem cache handle for vsk_anchor_t allocations */
94 kmem_cache_t *vsk_anchor_cache;

96 /* file events cleanup routine */
97 extern void free_fopdata(vnode_t *);

99 /*
100  * Root of AVL tree for the kstats associated with vopstats. Lock protects
101  * updates to vsktat_tree.
102  */
103 avl_tree_t      vskstat_tree;
104 kmutex_t        vskstat_tree_lock;

106 /* Global variable which enables/disables the vopstats collection */
107 int vopstats_enabled = 1;

109 /* Global used for empty/invalid v_path */
110 char *vn_vpath_empty = "";

112 /*
113  * forward declarations for internal vnode specific data (vsd)
114  */
115 static void *vsd_realloc(void *, size_t, size_t);

117 /*
118  * forward declarations for reparse point functions
119  */
120 static int fs_reparse_mark(char *target, vattn_t *vap, xvattr_t *xvattr);

122 /*
123  * VSD -- VNODE SPECIFIC DATA
124  * The v_data pointer is typically used by a file system to store a
125  * pointer to the file system's private node (e.g. ufs inode, nfs rnode).
126  * However, there are times when additional project private data needs

```

```

127 * to be stored separately from the data (node) pointed to by v_data.
128 * This additional data could be stored by the file system itself or
129 * by a completely different kernel entity. VSD provides a way for
130 * callers to obtain a key and store a pointer to private data associated
131 * with a vnode.
132 *
133 * Callers are responsible for protecting the vsd by holding v_vsd_lock
134 * for calls to vsd_set() and vsd_get().
135 */

137 /*
138 * vsd_lock protects:
139 *   vsd_nkeys - creation and deletion of vsd keys
140 *   vsd_list - insertion and deletion of vsd_node in the vsd_list
141 *   vsd_destructor - adding and removing destructors to the list
142 */
143 static kmutex_t      vsd_lock;
144 static uint_t        vsd_nkeys;      /* size of destructor array */
145 /* list of vsd_node's */
146 static list_t *vsd_list = NULL;
147 /* per-key destructor funcs */
148 static void          (**vsd_destructor)(void *);

150 /*
151 * The following is the common set of actions needed to update the
152 * vopstats structure from a vnode op. Both VOPSTATS_UPDATE() and
153 * VOPSTATS_UPDATE_IO() do almost the same thing, except for the
154 * recording of the bytes transferred. Since the code is similar
155 * but small, it is nearly a duplicate. Consequently any changes
156 * to one may need to be reflected in the other.
157 * Rundown of the variables:
158 * vp - Pointer to the vnode
159 * counter - Partial name structure member to update in vopstats for counts
160 * bytecounter - Partial name structure member to update in vopstats for bytes
161 * bytesval - Value to update in vopstats for bytes
162 * fstype - Index into vsanchor_fstype[], same as index into vfssw[]
163 * vsp - Pointer to vopstats structure (either in vfs or vsanchor_fstype[i])
164 */

166 #define VOPSTATS_UPDATE(vp, counter) {
167     vfs_t *vfsp = (vp)->v_vfsp;
168     if (vfsp && vfsp->vfs_impl &&
169         (vfsp->vfs_flag & VFS_STATS) && (vp)->v_type != VBAD) {
170         vopstats_t *vsp = &vfsp->vfs_vopstats;
171         uint64_t *stataddr = &(vsp->n##counter.value.ui64);
172         extern void __dtrace_probe__fsinfo_##counter(vnode_t *,
173             size_t, uint64_t *);
174         __dtrace_probe__fsinfo_##counter(vp, 0, stataddr);
175         (*stataddr)++;
176         if ((vsp = vfsp->vfs_fstypevsp) != NULL) {
177             vsp->n##counter.value.ui64++;
178         }
179     }
180 }

```

unchanged portion omitted

```

3315 /* VOP_XXX() macros call the corresponding fop_XXX() function */

3317 int
3318 fop_open(
3319     vnode_t **vpp,
3320     int mode,
3321     cred_t *cr,
3322     caller_context_t *ct)
3323 {
3324     int ret;

```

```

3325     vnode_t *vp = *vpp;

3327     VN_HOLD(vp);
3328     /*
3329     * Adding to the vnode counts before calling open
3330     * avoids the need for a mutex. It circumvents a race
3331     * condition where a query made on the vnode counts results in a
3332     * false negative. The inquirer goes away believing the file is
3333     * not open when there is an open on the file already under way.
3334     *
3335     * The counts are meant to prevent NFS from granting a delegation
3336     * when it would be dangerous to do so.
3337     *
3338     * The vnode counts are only kept on regular files
3339     */
3340     if ((*vpp)->v_type == VREG) {
3341         if (mode & FREAD)
3342             atomic_inc_32(&(*vpp)->v_rdcnt);
3343         if (mode & FWRITE)
3344             atomic_inc_32(&(*vpp)->v_wrcnt);
3345     }

3347     VOPXID_MAP_CR(vp, cr);

3349     ret = ((*vpp)->v_op->vop_open)(vpp, mode, cr, ct);

3351     if (ret) {
3352         /*
3353         * Use the saved vp just in case the vnode ptr got trashed
3354         * by the error.
3355         */
3356         VOPSTATS_UPDATE(vp, open);
3357         if ((vp->v_type == VREG) && (mode & FREAD))
3358             atomic_dec_32(&vp->v_rdcnt);
3359         if ((vp->v_type == VREG) && (mode & FWRITE))
3360             atomic_dec_32(&vp->v_wrcnt);
3361     } else {
3362         /*
3363         * Some filesystems will return a different vnode,
3364         * but the same path was still used to open it.
3365         * So if we do change the vnode and need to
3366         * copy over the path, do so here, rather than special
3367         * casing each filesystem. Adjust the vnode counts to
3368         * reflect the vnode switch.
3369         */
3370         VOPSTATS_UPDATE(*vpp, open);
3371         if (*vpp != vp) {
3372             if (*vpp != vp && *vpp != NULL) {
3373                 vn_copypath(vp, *vpp);
3374                 if (((*vpp)->v_type == VREG) && (mode & FREAD))
3375                     atomic_inc_32(&(*vpp)->v_rdcnt);
3376                 if ((vp->v_type == VREG) && (mode & FREAD))
3377                     atomic_dec_32(&vp->v_rdcnt);
3378                 if (((*vpp)->v_type == VREG) && (mode & FWRITE))
3379                     atomic_inc_32(&(*vpp)->v_wrcnt);
3380                 if ((vp->v_type == VREG) && (mode & FWRITE))
3381                     atomic_dec_32(&vp->v_wrcnt);
3382             }
3383             VN_RELE(vp);
3384             return (ret);
3385         }

```

unchanged portion omitted