

```

*****
116541 Mon Jan 14 13:17:03 2019
new/usr/src/uts/common/fs/vfs.c
10083 smatch fixes for common/fs/vfs.c
*****
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 2016 Joyent, Inc.
26  * Copyright 2016 Toomas Soome <tsoome@me.com>
27  * Copyright (c) 2016, 2017 by Delphix. All rights reserved.
28  * Copyright 2016 Nexenta Systems, Inc.
29  * Copyright 2017 RackTop Systems.
30 */

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

34 /*
35  * University Copyright- Copyright (c) 1982, 1986, 1988
36  * The Regents of the University of California
37  * All Rights Reserved
38  *
39  * University Acknowledgment- Portions of this document are derived from
40  * software developed by the University of California, Berkeley, and its
41  * contributors.
42  */

44 #include <sys/types.h>
45 #include <sys/t_lock.h>
46 #include <sys/param.h>
47 #include <sys/errno.h>
48 #include <sys/user.h>
49 #include <sys/fstyp.h>
50 #include <sys/kmem.h>
51 #include <sys/system.h>
52 #include <sys/proc.h>
53 #include <sys/mount.h>
54 #include <sys/vfs.h>
55 #include <sys/vfs_opreg.h>
56 #include <sys/fem.h>
57 #include <sys/mntent.h>
58 #include <sys/stat.h>
59 #include <sys/statvfs.h>
60 #include <sys/statfs.h>

```

```

61 #include <sys/cred.h>
62 #include <sys/vnode.h>
63 #include <sys/rwstlock.h>
64 #include <sys/dncl.h>
65 #include <sys/file.h>
66 #include <sys/time.h>
67 #include <sys/atomic.h>
68 #include <sys/cmn_err.h>
69 #include <sys/buf.h>
70 #include <sys/swap.h>
71 #include <sys/debug.h>
72 #include <sys/vnode.h>
73 #include <sys/modctl.h>
74 #include <sys/ddi.h>
75 #include <sys/pathname.h>
76 #include <sys/bootconf.h>
77 #include <sys/dumphdr.h>
78 #include <sys/dc_ki.h>
79 #include <sys/poll.h>
80 #include <sys/sunddi.h>
81 #include <sys/sysmacros.h>
82 #include <sys/zone.h>
83 #include <sys/policy.h>
84 #include <sys/ctfs.h>
85 #include <sys/objfs.h>
86 #include <sys/console.h>
87 #include <sys/reboot.h>
88 #include <sys/attr.h>
89 #include <sys/zio.h>
90 #include <sys/spa.h>
91 #include <sys/lofi.h>
92 #include <sys/bootprops.h>

94 #include <vm/page.h>

96 #include <fs/fs_subr.h>
97 /* Private interfaces to create vopstats-related data structures */
98 extern void initialize_vopstats(vopstats_t *);
99 extern vopstats_t *get_fstype_vopstats(struct vfs *, struct vfssw *);
100 extern vsk_anchor_t *get_vskstat_anchor(struct vfs *);

102 static void vfs_clearmntopt_nolock(mntopts_t *, const char *, int);
103 static void vfs_setmntopt_nolock(mntopts_t *, const char *,
104     const char *, int, int);
105 static int vfs_optionisset_nolock(const mntopts_t *, const char *, char **);
106 static void vfs_freemnttab(struct vfs *);
107 static void vfs_freeopt(mntopt_t *);
108 static void vfs_swapopttbl_nolock(mntopts_t *, mntopts_t *);
109 static void vfs_swapopttbl(mntopts_t *, mntopts_t *);
110 static void vfs_copyopttbl_extend(const mntopts_t *, mntopts_t *, int);
111 static void vfs_createopttbl_extend(mntopts_t *, const char *,
112     const mntopts_t *);
113 static char **vfs_copycancelopt_extend(char **const, int);
114 static void vfs_freecancelopt(char **);
115 static void getrootfs(char **, char **);
116 static int getmacpath(dev_info_t *, void *);
117 static void vfs_mnttabvvp_setup(void);

119 struct ipmnt {
120     struct ipmnt *mip_next;
121     dev_t mip_dev;
122     struct vfs *mip_vfsp;
123 };

unchanged_portion_omitted

1031 static void

```

```

1032 lofi_remove(struct vfs *vfsp)
1033 {
1034     struct lofi_ioctl *li = NULL;
1035     ldi_ident_t ldi_id;
1036     ldi_handle_t ldi_hdl;
1037     int err;

1039     if (vfsp->vfs_lofi_id == 0)
1040         return;

1042     ldi_id = ldi_ident_from_anon();

1044     li = kmem_zalloc(sizeof (*li), KM_SLEEP);
1045     li->li_id = vfs->vfs_lofi_id;
1046     li->li_cleanup = B_TRUE;

1048     err = ldi_open_by_name("/dev/lofictl", FREAD | FWRITE, kcred,
1049         &ldi_hdl, ldi_id);

1051     if (err)
1052         goto out;

1054     err = ldi_ioctl(ldi_hdl, LOFI_UNMAP_FILE_MINOR, (intptr_t)li,
1055         FREAD | FWRITE | FKIOCTL, kcred, NULL);

1057     (void) ldi_close(ldi_hdl, FREAD | FWRITE, kcred);

1059     if (!err)
1060         vfs->vfs_lofi_id = 0;

1062 out:
1063     ldi_ident_release(ldi_id);
1064     if (li != NULL)
1065         kmem_free(li, sizeof (*li));
1066 }
unchanged portion omitted

4157 vfs_t EIO_vfs;
4158 vfsops_t *EIO_vfsops;

4160 /*
4161  * Called from startup() to initialize all loaded vfs's
4162  */
4163 void
4164 vfsinit(void)
4165 {
4166     struct vfssw *vswp;
4167     int error;
4168     extern int vopstats_enabled;
4169     extern void vopstats_startup();

4171     static const fs_operation_def_t EIO_vfsops_template[] = {
4172         VFSNAME_MOUNT,          { .error = vfs_EIO },
4173         VFSNAME_UNMOUNT,       { .error = vfs_EIO },
4174         VFSNAME_ROOT,          { .error = vfs_EIO },
4175         VFSNAME_STATVFS,       { .error = vfs_EIO },
4176         VFSNAME_SYNC,          { .vfs_sync = vfs_EIO_sync },
4177         VFSNAME_VGET,          { .error = vfs_EIO },
4178         VFSNAME_MOUNTROOT,     { .error = vfs_EIO },
4179         VFSNAME_FREEVFS,       { .error = vfs_EIO },
4180         VFSNAME_VNSTATE,       { .error = vfs_EIO },
4181         NULL, NULL
4182     };

4184     static const fs_operation_def_t stray_vfsops_template[] = {
4185         VFSNAME_MOUNT,          { .error = vfststray },

```

```

4186         VFSNAME_UNMOUNT,       { .error = vfststray },
4187         VFSNAME_ROOT,          { .error = vfststray },
4188         VFSNAME_STATVFS,       { .error = vfststray },
4189         VFSNAME_SYNC,          { .vfs_sync = vfststray_sync },
4190         VFSNAME_VGET,          { .error = vfststray },
4191         VFSNAME_MOUNTROOT,     { .error = vfststray },
4192         VFSNAME_FREEVFS,       { .error = vfststray },
4193         VFSNAME_VNSTATE,       { .error = vfststray },
4194         NULL, NULL
4195     };

4197     /* Create vfs cache */
4198     vfs_cache = kmem_cache_create("vfs_cache", sizeof (struct vfs),
4199         sizeof (uintptr_t), NULL, NULL, NULL, NULL, 0);

4201     /* Initialize the vnode cache (file systems may use it during init). */
4202     vn_create_cache();

4204     /* Setup event monitor framework */
4205     fem_init();

4207     /* Initialize the dummy stray file system type. */
4208     error = vfs_setfsops(0, stray_vfsops_template, NULL);

4210     /* Initialize the dummy EIO file system. */
4211     error = vfs_makefsops(EIO_vfsops_template, &EIO_vfsops);
4212     if (error != 0) {
4213         cmn_err(CE_WARN, "vfsinit: bad EIO vfs ops template");
4214         /* Shouldn't happen, but not bad enough to panic */
4215     }

4217     VFS_INIT(&EIO_vfs, EIO_vfsops, (caddr_t)NULL);

4219     /*
4220      * Default EIO_vfs.vfs_flag to VFS_UNMOUNTED so a lookup
4221      * on this vfs can immediately notice it's invalid.
4222      */
4223     EIO_vfs.vfs_flag |= VFS_UNMOUNTED;

4225     /*
4226      * Call the init routines of non-loadable filesystems only.
4227      * Filesystems which are loaded as separate modules will be
4228      * initialized by the module loading code instead.
4229      */

4231     for (vswp = &vfssw[1]; vswp < &vfssw[nfstype]; vswp++) {
4232         RLOCK_VFSSW();
4233         if (vswp->vsw_init != NULL)
4234             (void) (*vswp->vsw_init)(vswp - vfssw, vswp->vsw_name);
4235         (*vswp->vsw_init)(vswp - vfssw, vswp->vsw_name);
4236         RUNLOCK_VFSSW();
4237     }

4238     vopstats_startup();

4240     if (vopstats_enabled) {
4241         /* EIO_vfs can collect stats, but we don't retrieve them */
4242         initialize_vopstats(&EIO_vfs.vfs_vopstats);
4243         EIO_vfs.vfs_fstypevswp = NULL;
4244         EIO_vfs.vfs_vskap = NULL;
4245         EIO_vfs.vfs_flag |= VFS_STATS;
4246     }

4248     xattr_init();

4250     reparse_point_init();

```

new/usr/src/uts/common/fs/vfs.c

5

4251 }  
unchanged\_portion\_omitted