

```

*****
17940 Tue Aug 31 13:02:02 2010
new/usr/src/uts/common/fs/smbclnt/smbfs/smbfs_client.c
6879933 Let SMBFS support extensible attributes per. PSARC 2007/315
*****
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) 2008, 2010, Oracle and/or its affiliates. All rights reserved.
24  *
25  * Copyright (c) 1983,1984,1985,1986,1987,1988,1989 AT&T.
26  * All rights reserved.
27 */

29 #include <sys/param.h>
30 #include <sys/system.h>
31 #include <sys/thread.h>
32 #include <sys/t_lock.h>
33 #include <sys/time.h>
34 #include <sys/vnode.h>
35 #include <sys/vfs.h>
36 #include <sys/errno.h>
37 #include <sys/buf.h>
38 #include <sys/stat.h>
39 #include <sys/cred.h>
40 #include <sys/kmem.h>
41 #include <sys/debug.h>
42 #include <sys/vmsystem.h>
43 #include <sys/flock.h>
44 #include <sys/share.h>
45 #include <sys/cmn_err.h>
46 #include <sys/tiuser.h>
47 #include <sys/sysmacros.h>
48 #include <sys/callb.h>
49 #include <sys/acl.h>
50 #include <sys/kstat.h>
51 #include <sys/signal.h>
52 #include <sys/list.h>
53 #include <sys/zone.h>

55 #include <netsmb/smb.h>
56 #include <netsmb/smb_conn.h>
57 #include <netsmb/smb_subr.h>

59 #include <smbfs/smbfs.h>
60 #include <smbfs/smbfs_node.h>
61 #include <smbfs/smbfs_subr.h>

```

```

63 #include <vm/hat.h>
64 #include <vm/as.h>
65 #include <vm/page.h>
66 #include <vm/pvn.h>
67 #include <vm/seg.h>
68 #include <vm/seg_map.h>
69 #include <vm/seg_vn.h>

71 static int smbfs_getattr_cache(vnode_t *, smbattr_t *);
72 static void smbattr_to_vattr(vnode_t *, smbattr_t *, vattr_t *);
73 static void smbattr_to_xvattr(vnode_t *, smbattr_t *, vattr_t *);
71 static int smbfs_getattr_cache(vnode_t *, struct smbattr *);
72 static int smbattr_to_vattr(vnode_t *, struct smbattr *,
73     struct vattr *);

75 /*
76  * The following code provide zone support in order to perform an action
77  * for each smbfs mount in a zone. This is also where we would add
78  * per-zone globals and kernel threads for the smbfs module (since
79  * they must be terminated by the shutdown callback).
80 */

82 struct smi_globals {
83     kmutex_t     smg_lock; /* lock protecting smg_list */
84     list_t       smg_list; /* list of SMBFS mounts in zone */
85     boolean_t    smg_destructor_called;
86 };
    unchanged_portion_omitted

404 /*
405  * Return either cached or remote attributes. If get remote attr
406  * use them to check and invalidate caches, then cache the new attributes.
407  *
408  * From NFS: nfsgetattr()
409  */
410 int
411 smbfsgetattr(vnode_t *vp, struct vattr *vap, cred_t *cr)
412 {
413     struct smbattr fa;
414     smbmntinfo_t *smi;
415     uint_t mask;
416     int error;

418     smi = VTOSMI(vp);

420     ASSERT(curproc->p_zone == smi->smi_zone_ref.zref_zone);

422     /*
423      * If asked for UID or GID, update n_uid, n_gid.
424      */
425     mask = AT_ALL;
426     if (vap->va_mask & (AT_UID | AT_GID)) {
427         if (smi->smi_flags & SMI_ACL)
428             (void) smbfs_acl_getids(vp, cr);
429         /* else leave as set in make_smbnode */
430     } else {
431         mask &= ~(AT_UID | AT_GID);
432     }

434     /*
435      * If we've got cached attributes, just use them;
436      * otherwise go to the server to get attributes,
437      * which will update the cache in the process.
438      */
439     error = smbfs_getattr_cache(vp, &fa);

```

```

440     if (error)
441         error = smbfs_getattr_otw(vp, &fa, cr);
442     if (error)
443         return (error);
444     vap->va_mask |= mask;
445 }
446 /*
447  * Re. client's view of the file size, see:
448  * smbfs_attrcache_fa, smbfs_getattr_otw
449  */
450     smbfattnr_to_vattr(vp, &fa, vap);
451     if (vap->va_mask & AT_XVATTR)
452         smbfattnr_to_xvattr(vp, &fa, vap);
453
454     return (0);
455     error = smbfattnr_to_vattr(vp, &fa, vap);
456     vap->va_mask = mask;
457
458     return (error);
459 }
460 /*
461  * Convert SMB over the wire attributes to vnode form.
462  * Returns 0 for success, error if failed (overflow, etc).
463  * From NFS: nattnr_to_vattr()
464  */
465 void
466 smbfattnr_to_vattr(vnode_t *vp, struct smbfattnr *fa, struct vattr *vap)
467 {
468     struct smbnode *np = VTOSMB(vp);
469
470     /* Set va_mask in caller */
471
472     /*
473      * Take type, mode, uid, gid from the smbfs node,
474      * which has have been updated by _getattr_otw.
475      */
476     vap->va_type = vp->v_type;
477     vap->va_mode = np->n_mode;
478
479     vap->va_uid = np->n_uid;
480     vap->va_gid = np->n_gid;
481
482     vap->va_fsid = vp->v_vfsp->vfs_dev;
483     vap->va_nodeid = np->n_ino;
484     vap->va_nlink = 1;
485
486     /*
487      * Difference from NFS here: We cache attributes as
488      * reported by the server, so r_attr.fa_size is the
489      * server's idea of the file size. This is called
490      * for getattr, so we want to return the client's
491      * idea of the file size. NFS deals with that in
492      * nfsgetattr(), the equivalent of our caller.
493      */
494     vap->va_size = np->r_size;
495
496     /*
497      * Times. Note, already converted from NT to
498      * Unix form (in the unmarshalling code).
499      */
500     vap->va_atime = fa->fa_atime;
501     vap->va_mtime = fa->fa_mtime;
502     vap->va_ctime = fa->fa_ctime;

```

```

503     /*
504      * rdev, blksize, seq are made up.
505      * va_nblocks is 512 byte blocks.
506      */
507     vap->va_rdev = vp->v_rdev;
508     vap->va_blksize = MAXBSIZE;
509     vap->va_nblocks = (fsblkcnt64_t)btod(np->r_attr.fa_allopsz);
510     vap->va_seq = 0;
511 }
512 /*
513  * smbfattnr_to_xvattr: like smbfattnr_to_vattr but for
514  * Extensible system attributes (PSARC 2007/315)
515  */
516 static void
517 smbfattnr_to_xvattr(vnode_t *vp, struct smbfattnr *fa, struct vattr *vap)
518 {
519     struct smbnode *np = VTOSMB(vp);
520     xvattr_t *xvap = (xvattr_t *)vap; /* *vap may be xvattr_t */
521     xoptattr_t *xoap = NULL;
522
523     if ((xoap = xva_getxoptattr(xvap)) == NULL)
524         return;
525
526     if (XVA_ISSET_REQ(xvap, XAT_CREATETIME)) {
527         xoap->xoa_createtime = fa->fa_createtime;
528         XVA_SET_RTN(xvap, XAT_CREATETIME);
529     }
530
531     if (XVA_ISSET_REQ(xvap, XAT_ARCHIVE)) {
532         xoap->xoa_archive =
533             ((fa->fa_attr & SMB_FA_ARCHIVE) != 0);
534         XVA_SET_RTN(xvap, XAT_ARCHIVE);
535     }
536
537     if (XVA_ISSET_REQ(xvap, XAT_SYSTEM)) {
538         xoap->xoa_system =
539             ((fa->fa_attr & SMB_FA_SYSTEM) != 0);
540         XVA_SET_RTN(xvap, XAT_SYSTEM);
541     }
542
543     if (XVA_ISSET_REQ(xvap, XAT_READONLY)) {
544         xoap->xoa_readonly =
545             ((fa->fa_attr & SMB_FA_RDONLY) != 0);
546         XVA_SET_RTN(xvap, XAT_READONLY);
547     }
548
549     if (XVA_ISSET_REQ(xvap, XAT_HIDDEN)) {
550         xoap->xoa_hidden =
551             ((fa->fa_attr & SMB_FA_HIDDEN) != 0);
552         XVA_SET_RTN(xvap, XAT_HIDDEN);
553     }
554     return (0);
555 }
556 /*
557  * SMB Client initialization and cleanup.
558  * Much of it is per-zone now.
559  */
560 /* ARGSUSED */
561 static void *
562 smbfs_zone_init(zoneid_t zoneid)

```

```
563 {
564     smi_globals_t *smg;

566     smg = kmem_alloc(sizeof (*smg), KM_SLEEP);
567     mutex_init(&smg->smg_lock, NULL, MUTEX_DEFAULT, NULL);
568     list_create(&smg->smg_list, sizeof (smbmntinfo_t),
569               offsetof(smbmntinfo_t, smi_zone_node));
570     smg->smg_destructor_called = E_FALSE;
571     return (smg);
572 }
unchanged_portion_omitted
```

```

*****
24755 Tue Aug 31 13:02:03 2010
new/usr/src/uts/common/fs/smbclnt/smbfs/smbfs_vfsops.c
6879933 Let SMBFS support extensible attributes per. PSARC 2007/315
*****
_____unchanged_portion_omitted_____

331 /*
332  * smbfs mount vfsop
333  * Set up mount info record and attach it to vfs struct.
334  */
335 static int
336 smbfs_mount(vfs_t *vfsp, vnode_t *mvp, struct mounta *uap, cred_t *cr)
337 {
338     char            *data = uap->dataptr;
339     int             error;
340     smbnode_t      *rtnp = NULL; /* root of this fs */
341     submntinfo_t   *smi = NULL;
342     dev_t          smbfs_dev;
343     int            version;
344     int            devfd;
345     zone_t         *zone = curproc->p_zone;
346     zone_t         *mntzone = NULL;
347     smb_share_t    *ssp = NULL;
348     smb_cred_t     scred;
349     int            flags, sec;

351     STRUCT_DECL(smbfs_args, args); /* smbfs mount arguments */

353     if ((error = secpolicy_fs_mount(cr, mvp, vfsp)) != 0)
354         return (error);

356     if (mvp->v_type != VDIR)
357         return (ENOTDIR);

359     /*
360      * get arguments
361      * uap->datalen might be different from sizeof (args)
362      * in a compatible situation.
363      */
364     STRUCT_INIT(args, get_udatamodel());
365     bzero(STRUCT_BUF(args), SIZEOF_STRUCT(smbfs_args, DATAMODEL_NATIVE));
366     if (copyin(data, STRUCT_BUF(args), MIN(uap->datalen,
367         SIZEOF_STRUCT(smbfs_args, DATAMODEL_NATIVE))))
368         return (EFAULT);
369

371     /*
372      * Check mount program version
373      */
374     version = STRUCT_FGET(args, version);
375     if (version != SMBFS_VERSION) {
376         cmn_err(CE_WARN, "mount version mismatch: "
377             " kernel=%d, mount=%d\n",
378             SMBFS_VERSION, version);
379         return (EINVAL);
380     }

382     /*
383      * Deal with re-mount requests.
384      */
385     if (uap->flags & MS_REMOUNT) {
386         cmn_err(CE_WARN, "MS_REMOUNT not implemented");
387         return (ENOTSUP);
388     }

```

```

390     /*
391      * Check for busy
392      */
393     mutex_enter(&mvp->v_lock);
394     if (!(uap->flags & MS_OVERLAY) &&
395         (mvp->v_count != 1 || (mvp->v_flag & VROOT))) {
396         mutex_exit(&mvp->v_lock);
397         return (EBUSY);
398     }
399     mutex_exit(&mvp->v_lock);

401     /*
402      * Get the "share" from the netsmb driver (ssp).
403      * It is returned with a "ref" (hold) for us.
404      * Release this hold: at errout below, or in
405      * smbfs_freevfs().
406      */
407     devfd = STRUCT_FGET(args, devfd);
408     error = smb_dev2share(devfd, &ssp);
409     if (error) {
410         cmn_err(CE_WARN, "invalid device handle %d (%d)\n",
411             devfd, error);
412         return (error);
413     }

415     /*
416      * Use "goto errout" from here on.
417      * See: ssp, smi, rtnp, mntzone
418      */

420     /*
421      * Determine the zone we're being mounted into.
422      */
423     zone_hold(mntzone = zone); /* start with this assumption */
424     if (getzoneid() == GLOBAL_ZONEID) {
425         zone_rele(mntzone);
426         mntzone = zone_find_by_path(refstr_value(vfsp->vfs_mntpt));
427         ASSERT(mntzone != NULL);
428         if (mntzone != zone) {
429             error = EBUSY;
430             goto errout;
431         }
432     }

434     /*
435      * Stop the mount from going any further if the zone is going away.
436      */
437     if (zone_status_get(mntzone) >= ZONE_IS_SHUTTING_DOWN) {
438         error = EBUSY;
439         goto errout;
440     }

442     /*
443      * On a Trusted Extensions client, we may have to force read-only
444      * for read-down mounts.
445      */
446     if (is_system_labeled()) {
447         void *addr;
448         int ipvers = 0;
449         struct smb_vc *vcp;

451         vcp = SSTOVC(ssp);
452         addr = smb_vc_getipaddr(vcp, &ipvers);
453         error = smbfs_mount_label_policy(vfsp, addr, ipvers, cr);

455         if (error > 0)

```

```

456         goto errout;
458         if (error == -1) {
459             /* change mount to read-only to prevent write-down */
460             vfs_setmntopt(vfsp, MNTOPT_RO, NULL, 0);
461         }
462     }
464     /* Prevent unload. */
465     atomic_inc_32(&smbfs_mountcount);
467     /*
468     * Create a mount record and link it to the vfs struct.
469     * No more possibilities for errors from here on.
470     * Tear-down of this stuff is in smbfs_free_smi()
471     *
472     * Compare with NFS: nfsrootvp()
473     */
474     smi = kmem_zalloc(sizeof (*smi), KM_SLEEP);
476     mutex_init(&smi->smi_lock, NULL, MUTEX_DEFAULT, NULL);
477     cv_init(&smi->smi_statvfs_cv, NULL, CV_DEFAULT, NULL);
479     rw_init(&smi->smi_hash_lk, NULL, RW_DEFAULT, NULL);
480     smbfs_init_hash_avl(&smi->smi_hash_avl);
482     smi->smi_share = ssp;
483     ssp = NULL;
485     /*
486     * Convert the anonymous zone hold acquired via zone_hold() above
487     * into a zone reference.
488     */
489     zone_init_ref(&smi->smi_zone_ref);
490     zone_hold_ref(mntzone, &smi->smi_zone_ref, ZONE_REF_SMBFS);
491     zone_rele(mntzone);
492     mntzone = NULL;
494     /*
495     * Initialize option defaults
496     */
497     smi->smi_flags = SMI_LLOCK;
498     smi->smi_acregmin = SEC2HR(SMBFS_ACREGMIN);
499     smi->smi_acregmax = SEC2HR(SMBFS_ACREGMAX);
500     smi->smi_acdirmin = SEC2HR(SMBFS_ACDIRMIN);
501     smi->smi_acdirmax = SEC2HR(SMBFS_ACDIRMAX);
503     /*
504     * All "generic" mount options have already been
505     * handled in vfs.c:domount() - see mntopts stuff.
506     * Query generic options using vfs_optionisset().
507     */
508     if (vfs_optionisset(vfsp, MNTOPT_INTR, NULL))
509         smi->smi_flags |= SMI_INT;
510     if (vfs_optionisset(vfsp, MNTOPT_ACL, NULL))
511         smi->smi_flags |= SMI_ACL;
513     /*
514     * Get the mount options that come in as smbfs_args,
515     * starting with args.flags (SMBFS_MF_XXX)
516     */
517     flags = STRUCT_FGET(args, flags);
518     smi->smi_uid = STRUCT_FGET(args, uid);
519     smi->smi_gid = STRUCT_FGET(args, gid);
520     smi->smi_fmode = STRUCT_FGET(args, file_mode) & 0777;
521     smi->smi_dmode = STRUCT_FGET(args, dir_mode) & 0777;

```

```

523     /*
524     * Handle the SMBFS_MF_XXX flags.
525     */
526     if (flags & SMBFS_MF_NOAC)
527         smi->smi_flags |= SMI_NOAC;
528     if (flags & SMBFS_MF_ACREGMIN) {
529         sec = STRUCT_FGET(args, acregmin);
530         if (sec < 0 || sec > SMBFS_ACMINMAX)
531             sec = SMBFS_ACMINMAX;
532         smi->smi_acregmin = SEC2HR(sec);
533     }
534     if (flags & SMBFS_MF_ACREGMAX) {
535         sec = STRUCT_FGET(args, acregmax);
536         if (sec < 0 || sec > SMBFS_ACMAXMAX)
537             sec = SMBFS_ACMAXMAX;
538         smi->smi_acregmax = SEC2HR(sec);
539     }
540     if (flags & SMBFS_MF_ACDIRMIN) {
541         sec = STRUCT_FGET(args, acdirmin);
542         if (sec < 0 || sec > SMBFS_ACMINMAX)
543             sec = SMBFS_ACMINMAX;
544         smi->smi_acdirmin = SEC2HR(sec);
545     }
546     if (flags & SMBFS_MF_ACDIRMAX) {
547         sec = STRUCT_FGET(args, acdirmax);
548         if (sec < 0 || sec > SMBFS_ACMAXMAX)
549             sec = SMBFS_ACMAXMAX;
550         smi->smi_acdirmax = SEC2HR(sec);
551     }
553     /*
554     * Get attributes of the remote file system,
555     * i.e. ACL support, named streams, etc.
556     */
557     smb_credinit(&scred, cr);
558     error = smbfs_smb_qfsattr(smi->smi_share, &smi->smi_fsa, &scred);
559     smb_credrele(&scred);
560     if (error) {
561         SMBVDEBUG("smbfs_smb_qfsattr error %d\n", error);
562     }
564     /*
565     * We enable XATTR by default (via smbfs_mntopts)
566     * but if the share does not support named streams,
567     * force the NOXATTR option (also clears XATTR).
568     * Caller will set or clear VFS_XATTR after this.
569     */
570     if ((smi->smi_fsattr & FILE_NAMED_STREAMS) == 0)
571         vfs_setmntopt(vfsp, MNTOPT_NOXATTR, NULL, 0);
573     /*
574     * Ditto ACLs (disable if not supported on this share)
575     */
576     if ((smi->smi_fsattr & FILE_PERSISTENT_ACLS) == 0) {
577         vfs_setmntopt(vfsp, MNTOPT_NOACL, NULL, 0);
578         smi->smi_flags &= ~SMI_ACL;
579     }
581     /*
582     * Assign a unique device id to the mount
583     */
584     mutex_enter(&smbfs_minor_lock);
585     do {
586         smbfs_minor = (smbfs_minor + 1) & MAXMIN32;
587         smbfs_dev = makedevice(smbfs_major, smbfs_minor);

```

```
588     } while (vfs_devismounted(smbfs_dev));
589     mutex_exit(&smbfs_minor_lock);

591     vfsp->vfs_dev = smbfs_dev;
592     vfs_make_fsid(&vfsp->vfs_fsid, smbfs_dev, smbfsfstyp);
593     vfsp->vfs_data = (caddr_t)smi;
594     vfsp->vfsfstype = smbfsfstyp;
595     vfsp->vfs_bsize = MAXBSIZE;
596     vfsp->vfs_bcount = 0;

598     smi->smi_vfsp = vfsp;
599     smbfs_zonelist_add(smi);      /* undo in smbfs_freevfs */

601     /* PSARC 2007/227 VFS Feature Registration */
602     vfs_set_feature(vfsp, VFSFT_XVATTR);
603     vfs_set_feature(vfsp, VFSFT_SYSATTR_VIEWS);

605     /*
606     * Create the root vnode, which we need in unmount
607     * for the call to smbfs_check_table(), etc.
608     * Release this hold in smbfs_unmount.
609     */
610     rtnp = smbfs_node_findcreate(smi, "\\\"", 1, NULL, 0, 0,
611     &smbfs_fattr0);
612     ASSERT(rtnp != NULL);
613     rtnp->r_vnode->v_type = VDIR;
614     rtnp->r_vnode->v_flag |= VROOT;
615     smi->smi_root = rtnp;

617     /*
618     * NFS does other stuff here too:
619     *   async worker threads
620     *   init kstats
621     *
622     * End of code from NFS nfsrootvp()
623     */
624     return (0);

626 errout:
627     vfsp->vfs_data = NULL;
628     if (smi != NULL)
629         smbfs_free_smi(smi);

631     if (mntzone != NULL)
632         zone_rele(mntzone);

634     if (ssp != NULL)
635         smb_share_rele(ssp);

637     return (error);
638 }
    unchanged_portion_omitted_
```

```

*****
78744 Tue Aug 31 13:02:04 2010
new/usr/src/uts/common/fs/smbclnt/smbfs/smbfs_vnops.c
6879933 Let SMBFS support extensible attributes per. PSARC 2007/315
*****
_____unchanged_portion_omitted_____

97 /*
98 * Turning this on causes nodes to be created in the cache
99 * during directory listings, normally avoiding a second
100 * OtW attribute fetch just after a readdir.
101 */
102 int smbfs_fastlookup = 1;

104 /* local static function defines */

106 static int smbfslookup_cache(vnode_t *, char *, int, vnode_t **,
107                               cred_t *);
108 static int smbfslookup(vnode_t *dvp, char *nm, vnode_t **vpp, cred_t *cr,
109                        int cache_ok, caller_context_t *);
110 static int smbfsrename(vnode_t *odvp, char *onm, vnode_t *ndvp, char *nnm,
111                       cred_t *cr, caller_context_t *);
112 static int smbfssetattr(vnode_t *, struct vattr *, int, cred_t *);
113 static int smbfs_accessx(void *, int, cred_t *);
114 static int smbfs_readdir(vnode_t *vp, uio_t *uio, cred_t *cr, int *eofp,
115                          caller_context_t *);
116 static void smbfs_rele_fid(smbnode_t *, struct smb_cred *);
117 static uint32_t xvattr_to_dosattr(smbnode_t *, struct vattr *);

119 /*
120 * These are the vnode ops routines which implement the vnode interface to
121 * the networked file system. These routines just take their parameters,
122 * make them look networkish by putting the right info into interface structs,
123 * and then calling the appropriate remote routine(s) to do the work.
124 *
125 * Note on directory name lookup cacheing: If we detect a stale fhandle,
126 * we purge the directory cache relative to that vnode. This way, the
127 * user won't get burned by the cache repeatedly. See <smbfs/smbnode.h> for
128 * more details on smbnode locking.
129 */

131 static int smbfs_open(vnode_t **, int, cred_t *, caller_context_t *);
132 static int smbfs_close(vnode_t *, int, int, offset_t, cred_t *,
133                        caller_context_t *);
134 static int smbfs_read(vnode_t *, struct uio *, int, cred_t *,
135                      caller_context_t *);
136 static int smbfs_write(vnode_t *, struct uio *, int, cred_t *,
137                       caller_context_t *);
138 static int smbfs_ioctl(vnode_t *, int, intp_t, int, cred_t *, int *,
139                       caller_context_t *);
140 static int smbfs_getattr(vnode_t *, struct vattr *, int, cred_t *,
141                          caller_context_t *);
142 static int smbfs_setattr(vnode_t *, struct vattr *, int, cred_t *,
143                          caller_context_t *);
144 static int smbfs_access(vnode_t *, int, int, cred_t *, caller_context_t *);
145 static int smbfs_fsync(vnode_t *, int, cred_t *, caller_context_t *);
146 static void smbfs_inactive(vnode_t *, cred_t *, caller_context_t *);
147 static int smbfs_lookup(vnode_t *, char *, vnode_t **, struct pathname *,
148                         int *, vnode_t *, cred_t *, caller_context_t *,
149                         int *, pathname_t *);
150 static int smbfs_create(vnode_t *, char *, struct vattr *, enum vcexcl,
151                        int, vnode_t **, cred_t *, int, caller_context_t *,
152                        vsecattr_t *);
153 static int smbfs_remove(vnode_t *, char *, cred_t *, caller_context_t *,
154                        int);
155 static int smbfs_rename(vnode_t *, char *, vnode_t *, char *, cred_t *,

```

```

156 caller_context_t *, int);
157 static int smbfs_mkdir(vnode_t *, char *, struct vattr *, vnode_t **,
158                        cred_t *, caller_context_t *, int, vsecattr_t *);
159 static int smbfs_rmdir(vnode_t *, char *, vnode_t *, cred_t *,
160                       caller_context_t *, int);
161 static int smbfs_readdir(vnode_t *, struct uio *, cred_t *, int *,
162                          caller_context_t *, int);
163 static int smbfs_rlock(vnode_t *, int, caller_context_t *);
164 static void smbfs_rwunlock(vnode_t *, int, caller_context_t *);
165 static int smbfs_seek(vnode_t *, offset_t, offset_t *, caller_context_t *);
166 static int smbfs_frlock(vnode_t *, int, struct flock64 *, int, offset_t,
167                        struct flk_callback *, cred_t *, caller_context_t *);
168 static int smbfs_space(vnode_t *, int, struct flock64 *, int, offset_t,
169                        cred_t *, caller_context_t *);
170 static int smbfs_pathconf(vnode_t *, int, ulong_t *, cred_t *,
171                            caller_context_t *);
172 static int smbfs_setsecattr(vnode_t *, vsecattr_t *, int, cred_t *,
173                              caller_context_t *);
174 static int smbfs_getsecattr(vnode_t *, vsecattr_t *, int, cred_t *,
175                              caller_context_t *);
176 static int smbfs_shrlock(vnode_t *, int, struct shrlock *, int, cred_t *,
177                           caller_context_t *);

179 /* Dummy function to use until correct function is ported in */
180 int noop_vnodeop() {
181     return (0);
182 }
_____unchanged_portion_omitted_____

839 /*
840 * Return either cached or remote attributes. If get remote attr
841 * use them to check and invalidate caches, then cache the new attributes.
842 *
843 * XXX
844 * This op should eventually support PSARC 2007/315, Extensible Attribute
845 * Interfaces, for richer metadata.
846 */
847 /* ARGSUSED */
848 static int
849 smbfs_getattr(vnode_t *vp, struct vattr *vap, int flags, cred_t *cr,
850              caller_context_t *ct)
851 {
852     smbnode_t *np;
853     smbmntinfo_t *smi;
854
855     smi = VTOSMI(vp);
856
857     if (curproc->p_zone != smi->smi_zone_ref.zref_zone)
858         return (EIO);
859
860     if (smi->smi_flags & SMI_DEAD || vp->v_vfsp->vifs_flag & VFS_UNMOUNTED)
861         return (EIO);
862
863     /*
864      * If it has been specified that the return value will
865      * just be used as a hint, and we are only being asked
866      * for size, fsid or rdevid, then return the client's
867      * notion of these values without checking to make sure
868      * that the attribute cache is up to date.
869      * The whole point is to avoid an over the wire GETATTR
870      * call.
871      */
872     np = VTOSMB(vp);
873     if (flags & ATTR_HINT) {
874         if (vap->va_mask ==

```

```

871         (vap->va_mask & (AT_SIZE | AT_FSID | AT_RDEV))) {
872             mutex_enter(&np->r_statelock);
873             if (vap->va_mask | AT_SIZE)
874                 vap->va_size = np->r_size;
875             if (vap->va_mask | AT_FSID)
876                 vap->va_fsid = vp->v_vfsp->vfs_dev;
877             if (vap->va_mask | AT_RDEV)
878                 vap->va_rdev = vp->v_rdev;
879             mutex_exit(&np->r_statelock);
880             return (0);
881         }
882     }
883
884     return (smbfsgetattr(vp, vap, cr));
885 }
886
887 /* smbfssetattr() in smbfs_client.c */
888
889 /*
890  * XXX
891  * This op should eventually support PSARC 2007/315, Extensible Attribute
892  * Interfaces, for richer metadata.
893  */
894 /* ARGSUSED4 */
895 static int
896 smbfs_setattr(vnode_t *vp, struct vattn *vap, int flags, cred_t *cr,
897              caller_context_t *ct)
898 {
899     vfs_t          *vfsp;
900     smbmntinfo_t   *smi;
901     int             error;
902     uint_t         mask;
903     struct vattn   oldva;
904
905     vfsp = vp->v_vfsp;
906     smi = VFTOSMI(vfsp);
907
908     if (curproc->p_zone != smi->smi_zone_ref.zref_zone)
909         return (EIO);
910
911     if (smi->smi_flags & SMI_DEAD || vfsp->vfs_flag & VFS_UNMOUNTED)
912         return (EIO);
913
914     mask = vap->va_mask;
915     if (mask & AT_NOSET)
916         return (EINVAL);
917
918     if (vfsp->vfs_flag & VFS_RDONLY)
919         return (EROFS);
920
921     /*
922      * This is a _local_ access check so that only the owner of
923      * this mount can set attributes. With ACLs enabled, the
924      * file owner can be different from the mount owner, and we
925      * need to check the _mount_ owner here. See _access_rwx
926      */
927     bzero(&oldva, sizeof (oldva));
928     oldva.va_mask = AT_TYPE | AT_MODE;
929     error = smbfsgetattr(vp, &oldva, cr);
930     if (error)
931         return (error);
932     oldva.va_mask |= AT_UID | AT_GID;
933     oldva.va_uid = smi->smi_uid;
934     oldva.va_gid = smi->smi_gid;
935
936     error = secpolicy_vnode_setattr(cr, vp, vap, &oldva, flags,

```

```

932         smbfs_accessx, vp);
933     if (error)
934         return (error);
935
936     if (mask & (AT_UID | AT_GID)) {
937         if (smi->smi_flags & SMI_ACL)
938             error = smbfs_acl_setids(vp, vap, cr);
939         else
940             error = ENOSYS;
941         if (error != 0) {
942             SMBVDEBUG("error %d setting UID/GID on %s",
943                     error, VTOSMB(vp)->n_rpath);
944             /*
945              * It might be more correct to return the
946              * error here, but that causes complaints
947              * when root extracts a cpio archive, etc.
948              * So ignore this error, and go ahead with
949              * the rest of the setattr work.
950              */
951         }
952     }
953
954     return (smbfssetattr(vp, vap, flags, cr));
955 }
956
957 /*
958  * Mostly from Darwin smbfs_setattr()
959  * but then modified a lot.
960  */
961 /* ARGSUSED */
962 static int
963 smbfssetattr(vnode_t *vp, struct vattn *vap, int flags, cred_t *cr)
964 {
965     int             error = 0;
966     smbnode_t      *np = VTOSMB(vp);
967     uint_t         mask = vap->va_mask;
968     struct timespec *mtime, *atime;
969     struct smb_cred scred;
970     int            cerror, modified = 0;
971     unsigned short fid;
972     int have_fid = 0;
973     uint32_t rights = 0;
974     uint32_t dosattr = 0;
975
976     ASSERT(curproc->p_zone == VTOSMI(vp)->smi_zone_ref.zref_zone);
977
978     /*
979      * There are no settable attributes on the XATTR dir,
980      * so just silently ignore these. On XATTR files,
981      * you can set the size but nothing else.
982      */
983     if (vp->v_flag & V_XATTRDIR)
984         return (0);
985     if (np->n_flag & N_XATTR) {
986         if (mask & AT_TIMES)
987             SMBVDEBUG("ignore set time on xattr\n");
988         mask &= AT_SIZE;
989     }
990
991     /*
992      * If our caller is trying to set multiple attributes, they
993      * can make no assumption about what order they are done in.
994      * Here we try to do them in order of decreasing likelihood
995      * of failure, just to minimize the chance we'll wind up
996      * with a partially complete request.
997      */

```



```

999      /* Shared lock for (possible) n_fid use. */
1000     if (smbfs_rw_enter_sig(&np->r_lkserlock, RW_READER, SMBINTR(vp)))
1001         return (EINTR);
1002     smb_credinit(&scred, cr);

1004     /*
1005      * If the caller has provided extensible attributes,
1006      * map those into DOS attributes supported by SMB.
1007      * Note: zero means "no change".
1008      */
1009     if (mask & AT_XVATTR)
1010         dosattr = xvattr_to_dosattr(np, vap);

1012     /*
1013      * Will we need an open handle for this setattr?
1014      * If so, what rights will we need?
1015      */
1016     if (dosattr || (mask & (AT_ETIME | AT_MTIME))) {
1017         if (mask & (AT_ETIME | AT_MTIME)) {
1018             rights |=
1019                 SA_RIGHT_FILE_WRITE_ATTRIBUTES;
1020         }
1021         if (mask & AT_SIZE) {
1022             rights |=
1023                 SA_RIGHT_FILE_WRITE_DATA |
1024                 SA_RIGHT_FILE_APPEND_DATA;
1025         }

1026         /*
1027          * Only SIZE really requires a handle, but it's
1028          * simpler and more reliable to set via a handle.
1029          * Some servers like NT4 won't set times by path.
1030          * Also, we're usually setting everything anyway.
1031          */
1032         if (rights != 0) {
1033             if (mask & (AT_SIZE | AT_ETIME | AT_MTIME)) {
1034                 error = smbfs_smb_tmpopen(np, rights, &scred, &fid);
1035                 if (error) {
1036                     SMBVDEBUG("error %d opening %s\n",
1037                             error, np->n_rpath);
1038                     goto out;
1039                 }
1040                 have_fid = 1;
1041             }

1042             /*
1043              * If the server supports the UNIX extensions, right here is where
1044              * we'd support changes to uid, gid, mode, and possibly va_flags.
1045              * For now we claim to have made any such changes.
1046              */

1048             if (mask & AT_SIZE) {
1049                 /*
1050                  * If the new file size is less than what the client sees as
1051                  * the file size, then just change the size and invalidate
1052                  * the pages.
1053                  * I am commenting this code at present because the function
1054                  * smbfs_putapage() is not yet implemented.
1055                  */

1057                 /*
1058                  * Set the file size to vap->va_size.
1059                  */
1060                 ASSERT(have_fid);
1061                 error = smbfs_smb_setfsize(np, fid, vap->va_size, &scred);

```

```

1062         if (error) {
1063             SMBVDEBUG("setsize error %d file %s\n",
1064                     error, np->n_rpath);
1065         } else {
1066             /*
1067              * Darwin had code here to zero-extend.
1068              * Tests indicate the server will zero-fill,
1069              * so looks like we don't need to do this.
1070              * Good thing, as this could take forever.
1071              *
1072              * XXX: Reportedly, writing one byte of zero
1073              * at the end offset avoids problems here.
1074              */
1075             mutex_enter(&np->r_stalock);
1076             np->r_size = vap->va_size;
1077             mutex_exit(&np->r_stalock);
1078             modified = 1;
1079         }
1080     }

1082     /*
1083      * XXX: When Solaris has create_time, set that too.
1084      * Note: create_time is different from ctime.
1085      */
1086     mtime = ((mask & AT_MTIME) ? &vap->va_mtime : 0);
1087     atime = ((mask & AT_ETIME) ? &vap->va_atime : 0);

1089     if (dosattr || mtime || atime) {
1090         if (mtime || atime) {
1091             /*
1092              * Always use the handle-based set attr call now.
1093              * Not trying to set DOS attributes here so pass zero.
1094              */
1095             ASSERT(have_fid);
1096             error = smbfs_smb_setfattr(np, fid,
1097                                     dosattr, mtime, atime, &scred);
1098             if (error) {
1099                 SMBVDEBUG("set times error %d file %s\n",
1100                         error, np->n_rpath);
1101             } else {
1102                 modified = 1;
1103             }
1104         }

1105     out:
1106         if (modified) {
1107             /*
1108              * Invalidate attribute cache in case the server
1109              * doesn't set exactly the attributes we asked.
1110              */
1111             smbfs_attrcache_remove(np);
1112         }

1114         if (have_fid) {
1115             cerr = smbfs_smb_tmpclose(np, fid, &scred);
1116             if (cerr)
1117                 SMBVDEBUG("error %d closing %s\n",
1118                         cerr, np->n_rpath);
1119         }

1121         smb_credrele(&scred);
1122         smbfs_rw_exit(&np->r_lkserlock);

1124         return (error);
1125     }

```

```

1127 /*
1128 * Helper function for extensible system attributes (PSARC 2007/315)
1129 * Compute the DOS attribute word to pass to _setfatattr (see above).
1130 * This returns zero IFF no change is being made to attributes.
1131 * Otherwise return the new attributes or SMB_EFA_NORMAL.
1132 */
1133 static uint32_t
1134 xvattr_to_dosattr(smbnode_t *np, struct vattr *vap)
1135 {
1136     xvattr_t *xvap = (xvattr_t *)vap;
1137     xoattr_t *xoap = NULL;
1138     uint32_t attr = np->r_attr.fa_attr;
1139     boolean_t anyset = B_FALSE;
1140
1141     if ((xoap = xva_getxoattr(xvap)) == NULL)
1142         return (0);
1143
1144     if (XVA_ISSET_REQ(xvap, XAT_ARCHIVE)) {
1145         if (xoap->xoa_archive)
1146             attr |= SMB_FA_ARCHIVE;
1147         else
1148             attr &= ~SMB_FA_ARCHIVE;
1149         XVA_SET_RTN(xvap, XAT_ARCHIVE);
1150         anyset = B_TRUE;
1151     }
1152     if (XVA_ISSET_REQ(xvap, XAT_SYSTEM)) {
1153         if (xoap->xoa_system)
1154             attr |= SMB_FA_SYSTEM;
1155         else
1156             attr &= ~SMB_FA_SYSTEM;
1157         XVA_SET_RTN(xvap, XAT_SYSTEM);
1158         anyset = B_TRUE;
1159     }
1160     if (XVA_ISSET_REQ(xvap, XAT_READONLY)) {
1161         if (xoap->xoa_readonly)
1162             attr |= SMB_FA_RDONLY;
1163         else
1164             attr &= ~SMB_FA_RDONLY;
1165         XVA_SET_RTN(xvap, XAT_READONLY);
1166         anyset = B_TRUE;
1167     }
1168     if (XVA_ISSET_REQ(xvap, XAT_HIDDEN)) {
1169         if (xoap->xoa_hidden)
1170             attr |= SMB_FA_HIDDEN;
1171         else
1172             attr &= ~SMB_FA_HIDDEN;
1173         XVA_SET_RTN(xvap, XAT_HIDDEN);
1174         anyset = B_TRUE;
1175     }
1176
1177     if (anyset == B_FALSE)
1178         return (0); /* no change */
1179     if (attr == 0)
1180         attr = SMB_EFA_NORMAL;
1181
1182     return (attr);
1183 }
1184
1185 /*
1186 * smbfs_access_rwx()
1187 * Common function for smbfs_access, etc.
1188 *
1189 * The security model implemented by the FS is unusual
1190 * due to the current "single user mounts" restriction:
1191 * All access under a given mount point uses the CIFS

```

```

1192 * credentials established by the owner of the mount.
1193 *
1194 * Most access checking is handled by the CIFS server,
1195 * but we need sufficient Unix access checks here to
1196 * prevent other local Unix users from having access
1197 * to objects under this mount that the uid/gid/mode
1198 * settings in the mount would not allow.
1199 *
1200 * With this model, there is a case where we need the
1201 * ability to do an access check before we have the
1202 * vnode for an object. This function takes advantage
1203 * of the fact that the uid/gid/mode is per mount, and
1204 * avoids the need for a vnode.
1205 *
1206 * We still (sort of) need a vnode when we call
1207 * secpolicy_vnode_access, but that only uses
1208 * the vtype field, so we can use a pair of fake
1209 * vnodes that have only v_type filled in.
1210 *
1211 * XXX: Later, add a new secpolicy_vtype_access()
1212 * that takes the vtype instead of a vnode, and
1213 * get rid of the tmpl_vxxx fake vnodes below.
1214 */
1215 static int
1216 smbfs_access_rwx(vfs_t *vfsp, int vtype, int mode, cred_t *cr)
1217 {
1218     /* See the secpolicy call below. */
1219     static const vnode_t tmpl_vdir = { .v_type = VDIR };
1220     static const vnode_t tmpl_vreg = { .v_type = VREG };
1221     vattr_t va;
1222     vnode_t *tvp;
1223     struct smbmntinfo *smi = VFTOSMI(vfsp);
1224     int shift = 0;
1225
1226     /*
1227      * Build our (fabricated) vnode attributes.
1228      * XXX: Could make these templates in the
1229      * per-mount struct and use them here.
1230      */
1231     bzero(&va, sizeof (va));
1232     va.va_mask = AT_TYPE | AT_MODE | AT_UID | AT_GID;
1233     va.va_type = vtype;
1234     va.va_mode = (vtype == VDIR) ?
1235         smi->smi_dmode : smi->smi_fmmode;
1236     va.va_uid = smi->smi_uid;
1237     va.va_gid = smi->smi_gid;
1238
1239     /*
1240      * Disallow write attempts on read-only file systems,
1241      * unless the file is a device or fifo node. Note:
1242      * Inline vn_is_readonly and IS_DEVVP here because
1243      * we may not have a vnode ptr. Original expr. was:
1244      * (mode & VWRITE) && vn_is_readonly(vp) && !IS_DEVVP(vp)
1245      */
1246     if ((mode & VWRITE) &&
1247         (vfsp->vfs_flag & VFS_RDONLY) &&
1248         !(vtype == VCHR || vtype == VBLK || vtype == VFIFO))
1249         return (EROFS);
1250
1251     /*
1252      * Disallow attempts to access mandatory lock files.
1253      * Similarly, expand MANDLOCK here.
1254      * XXX: not sure we need this.
1255      */
1256     if ((mode & (VWRITE | VREAD | VEXEC)) &&
1257         va.va_type == VREG && MANDMODE(va.va_mode))

```

```

1258         return (EACCES);

1260     /*
1261     * Access check is based on only
1262     * one of owner, group, public.
1263     * If not owner, then check group.
1264     * If not a member of the group,
1265     * then check public access.
1266     */
1267     if (crgetuid(cr) != va.va_uid) {
1268         shift += 3;
1269         if (!groupmember(va.va_gid, cr))
1270             shift += 3;
1271     }

1273     /*
1274     * We need a vnode for secpolicy_vnode_access,
1275     * but the only thing it looks at is v_type,
1276     * so pass one of the templates above.
1277     */
1278     tvp = (va.va_type == VDIR) ?
1279         (vnode_t *)&tmpl_vdir :
1280         (vnode_t *)&tmpl_vreg;

1282     return (secpolicy_vnode_access2(cr, tvp, va.va_uid,
1283         va.va_mode << shift, mode));
1284 }

```

unchanged portion omitted

```

3001 /* ARGSUSED */
3002 static int
3003 smbfs_pathconf(vnode_t *vp, int cmd, ulong_t *valp, cred_t *cr,
3004     caller_context_t *ct)
3005 {
3006     vfs_t *vfs;
3007     smbmntinfo_t *smi;
3008     struct smb_share *ssp;

3010     vfs = vp->v_vfsp;
3011     smi = VFTOSMI(vfs);

3013     if (curproc->p_zone != smi->smi_zone_ref.zref_zone)
3014         return (EIO);

3016     if (smi->smi_flags & SMI_DEAD || vp->v_vfsp->vfs_flag & VFS_UNMOUNTED)
3017         return (EIO);

3019     switch (cmd) {
3020     case _PC_FILESIZEBITS:
3021         ssp = smi->smi_share;
3022         if (SSTOVC(ssp)->vc_sopt.sv_caps & SMB_CAP_LARGE_FILES)
3023             *valp = 64;
3024         else
3025             *valp = 32;
3026         break;

3028     case _PC_LINK_MAX:
3029         /* We only ever report one link to an object */
3030         *valp = 1;
3031         break;

3033     case _PC_ACL_ENABLED:
3034         /*
3035         * Always indicate that ACLs are enabled and
3036         * that we support ACE_T format, otherwise
3037         * libsec will ask for ACLENT_T format data

```

```

3038         * which we don't support.
3039         */
3040         *valp = _ACL_ACE_ENABLED;
3041         break;

3043     case _PC_SYMLINK_MAX: /* No symlinks until we do Unix extensions */
3044         *valp = 0;
3045         break;

3047     case _PC_XATTR_EXISTS:
3048         if (vfs->vfs_flag & VFS_XATTR) {
3049             *valp = smbfs_xa_exists(vp, cr);
3050             break;
3051         }
3052         return (EINVAL);

3054     case _PC_SATTR_ENABLED:
3055     case _PC_SATTR_EXISTS:
3056         *valp = 1;
3057         break;

3059     case _PC_TIMESTAMP_RESOLUTION:
3060         /*
3061         * Windows times are tenths of microseconds
3062         * (multiples of 100 nanoseconds).
3063         */
3064         *valp = 100L;
3065         break;

3067     default:
3068         return (fs_pathconf(vp, cmd, valp, cr, ct));
3069     }
3070     return (0);
3071 }

```

unchanged portion omitted

```

*****
39409 Tue Aug 31 13:02:05 2010
new/usr/src/uts/common/fs/xattr.c
6975745 xattr directories should be more transparent
*****
_____unchanged_portion_omitted_____

55 typedef struct {
56     gfs_dir_t      xattr_gfs_private;
57     vnode_t        *xattr_realvp;
57     vnode_t        *xattr_realvp; /* Only used for VOP_REALVP */
58 } xattr_dir_t;

60 /*
61  * xattr_realvp is only used for VOP_REALVP, this is so we don't
62  * keep an unnecessary hold on the *real* xattr dir unless we have
63  * no other choice.
64  */

60 /* ARGSUSED */
61 static int
62 xattr_file_open(vnode_t **vpp, int flags, cred_t *cr, caller_context_t *ct)
63 {
64     xattr_file_t *np = (*vpp)->v_data;

66     if ((np->xattr_view == XATTR_VIEW_READONLY) && (flags & FWRITE))
67         return (EACCES);

69     return (0);
70 }
_____unchanged_portion_omitted_____

872 static int
873 xattr_dir_readdir(vnode_t *dvp, vnode_t **realvp, int lookup_flags,
874     cred_t *cr, caller_context_t *ct)
875 {
876     xattr_dir_t *xattr_dir;
882     vnode_t *pvp;
877     int error;
884     struct pathname pn;
885     char *startnm = "";

879     *realvp = NULL;

881     if (dvp->v_type != VDIR)
882         return (EINVAL);
889     pvp = gfs_file_parent(dvp);

884     mutex_enter(&dvp->v_lock);
885     xattr_dir = dvp->v_data;
886     *realvp = xattr_dir->xattr_realvp;
887     mutex_exit(&dvp->v_lock);
891     error = pn_get(startnm, UIO_SYSSPACE, &pn);
892     if (error) {
893         VN_RELE(pvp);
894         return (error);
895     }

889     if (*realvp != NULL) {
890         VN_HOLD(*realvp);
891         error = 0;
892     } else
893         error = ENOENT;
897     /*
898     * Set the LOOKUP_HAVE_SYSATTR_DIR flag so that we don't get into an
899     * infinite loop with fop_lookup calling back to xattr_dir_lookup.

```

```

900     */
901     lookup_flags |= LOOKUP_HAVE_SYSATTR_DIR;
902     error = VOP_LOOKUP(pvp, startnm, realvp, &pn, lookup_flags,
903         rootvp, cr, ct, NULL, NULL);
904     pn_free(&pn);

895     return (error);
896 }

898 /* ARGSUSED */
899 static int
900 xattr_dir_open(vnode_t **vpp, int flags, cred_t *cr, caller_context_t *ct)
901 {
902     vnode_t *realvp;
903     int error;

905     if (flags & FWRITE) {
906         return (EACCES);
907     }

909     /*
910     * The underlying FS may need this VOP call.
911     */
912     error = xattr_dir_readdir(*vpp, &realvp, LOOKUP_XATTR, cr, ct);
913     if (error == 0) {
914         error = VOP_OPEN(&realvp, flags, cr, ct);
915         VN_RELE(realvp);
916         if (error)
917             return (error);
918     } /* else ignore this error */

920     return (0);
921 }

923 /* ARGSUSED */
924 static int
925 xattr_dir_close(vnode_t *vp, int flags, int count, offset_t off, cred_t *cr,
926     xattr_dir_close(vnode_t *vpp, int flags, int count, offset_t off, cred_t *cr,
927     caller_context_t *ct)
928 {
929     vnode_t *realvp;
930     int error;

931     /*
932     * The underlying FS may need this VOP call.
933     */
934     error = xattr_dir_readdir(vp, &realvp, LOOKUP_XATTR, cr, ct);
935     if (error == 0) {
936         error = VOP_CLOSE(realvp, flags, count, off, cr, ct);
937         VN_RELE(realvp);
938         if (error)
939             return (error);
940     } /* else ignore this error */

942     return (0);
943 }
_____unchanged_portion_omitted_____

1372 /* ARGSUSED */
1373 static int
1374 xattr_dir_realvp(vnode_t *vp, vnode_t **realvp, caller_context_t *ct)
1375 {
1359     xattr_dir_t *xattr_dir;

1361     mutex_enter(&vp->v_lock);
1362     xattr_dir = vp->v_data;

```

```

1363     if (xattr_dir->xattr_realvp) {
1364         *realvp = xattr_dir->xattr_realvp;
1365         mutex_exit(&vp->v_lock);
1366         return (0);
1367     } else {
1368         vnode_t *xdvp;
1376     int error;

1378     error = xattr_dir_readdir(vp, realvp, LOOKUP_XATTR, kcred, NULL);
1379     mutex_exit(&vp->v_lock);
1380     if ((error = xattr_dir_readdir(vp, &xdvp,
1381         LOOKUP_XATTR, kcred, NULL)) == 0) {
1382         /*
1383          * verify we aren't racing with another thread
1384          * to find the xattr_realvp
1385          */
1386         mutex_enter(&vp->v_lock);
1387         if (xattr_dir->xattr_realvp == NULL) {
1388             xattr_dir->xattr_realvp = xdvp;
1389             *realvp = xdvp;
1390             mutex_exit(&vp->v_lock);
1391         } else {
1392             *realvp = xattr_dir->xattr_realvp;
1393             mutex_exit(&vp->v_lock);
1394             VN_RELE(xdvp);
1395         }
1396     }
1397     return (error);
1398 }
1399 }
1400 }
1401 }
1402 }
1403 }
1404 }
1405 }
1406 }
1407 }
1408 }
1409 }
1410 }
1411 }
1412 }
1413 }
1414 }
1415 }
1416 }
1417 }
1418 }
1419 }
1420 }
1421 }
1422 }
1423 }
1424 }
1425 }
1426 }
1427 }
1428 }
1429 }
1430 }
1431 }
1432 }
1433 }
1434 }
1435 }
1436 }
1437 }
1438 }
1439 }
1440 }
1441 }
1442 }
1443 }
1444 }
1445 }
1446 }
1447 }
1448 }
1449 }
1450 }
1451 }
1452 }
1453 }
1454 }
1455 }
1456 }
1457 }
1458 }
1459 }
1460 }
1461 }
1462 }
1463 }
1464 }
1465 }
1466 }
1467 }
1468 }
1469 }
1470 }
1471 }
1472 }
1473 }
1474 }
1475 }
1476 }
1477 }
1478 }
1479 }
1480 }
1481 }
1482 }
1483 }
1484 }
1485 }
1486 }
1487 }
1488 }
1489 }
1490 }
1491 }
1492 }
1493 }
1494 }
1495 }
1496 }
1497 }
1498 }
1499 }
1500 }
1501 }
1502 }
1503 }
1504 }
1505 }
1506 }
1507 }
1508 }
1509 }
1510 }
1511 }
1512 }
1513 }
1514 }
1515 }
1516 }
1517 }
1518 }
1519 }
1520 }
1521 }
1522 }
1523 }
1524 }
1525 }
1526 }
1527 }
1528 }
1529 }
1530 }
1531 }
1532 }
1533 }
1534 }
1535 }
1536 }
1537 }
1538 }
1539 }
1540 }
1541 }
1542 }
1543 }
1544 }
1545 }
1546 }
1547 }
1548 }
1549 }
1550 }
1551 }
1552 }
1553 }
1554 }
1555 }
1556 }
1557 }
1558 }
1559 }
1560 }
1561 }
1562 }
1563 }
1564 }
1565 }
1566 }
1567 }
1568 }
1569 }
1570 }
1571 }
1572 }
1573 }
1574 }
1575 }
1576 }
1577 }
1578 }
1579 }
1580 }
1581 }
1582 }
1583 }
1584 }
1585 }
1586 }
1587 }
1588 }
1589 }
1590 }
1591 }
1592 }
1593 }
1594 }
1595 }
1596 }
1597 }
1598 }
1599 }
1600 }
1601 }
1602 }
1603 }
1604 }
1605 }
1606 }
1607 }
1608 }
1609 }
1610 }
1611 }
1612 }
1613 }
1614 }
1615 }
1616 }
1617 }
1618 }
1619 }
1620 }
1621 }
1622 }
1623 }
1624 }
1625 }
1626 }
1627 }
1628 }
1629 }
1630 }
1631 }
1632 }
1633 }
1634 }
1635 }
1636 }
1637 }
1638 }
1639 }
1640 }
1641 }
1642 }
1643 }
1644 }
1645 }
1646 }
1647 }
1648 }
1649 }
1650 }
1651 }
1652 }
1653 }
1654 }
1655 }
1656 }
1657 }
1658 }
1659 }
1660 }
1661 }
1662 }
1663 }
1664 }
1665 }
1666 }
1667 }
1668 }
1669 }
1670 }
1671 }
1672 }
1673 }
1674 }
1675 }
1676 }
1677 }
1678 }
1679 }
1680 }
1681 }
1682 }
1683 }
1684 }
1685 }
1686 }
1687 }
1688 }
1689 }
1690 }
1691 }
1692 }
1693 }
1694 }
1695 }
1696 }
1697 }
1698 }
1699 }
1700 }
1701 }
1702 }
1703 }
1704 }
1705 }
1706 }
1707 }
1708 }
1709 }
1710 }
1711 }
1712 }
1713 }
1714 }
1715 }
1716 }
1717 }
1718 }
1719 }
1720 }
1721 }
1722 }
1723 }
1724 }
1725 }
1726 }
1727 }
1728 }
1729 }
1730 }
1731 }
1732 }
1733 }
1734 }
1735 }
1736 }
1737 }
1738 }
1739 }
1740 }
1741 }
1742 }
1743 }
1744 }
1745 }
1746 }
1747 }
1748 }
1749 }
1750 }
1751 }
1752 }
1753 }
1754 }
1755 }
1756 }
1757 }
1758 }
1759 }
1760 }
1761 }
1762 }
1763 }
1764 }
1765 }
1766 }
1767 }
1768 }
1769 }
1770 }
1771 }
1772 }
1773 }
1774 }
1775 }
1776 }
1777 }
1778 }
1779 }
1780 }
1781 }
1782 }
1783 }
1784 }
1785 }
1786 }
1787 }
1788 }
1789 }
1790 }
1791 }
1792 }
1793 }
1794 }
1795 }
1796 }
1797 }
1798 }
1799 }
1800 }
1801 }
1802 }
1803 }
1804 }
1805 }
1806 }
1807 }
1808 }
1809 }
1810 }
1811 }
1812 }
1813 }
1814 }
1815 }
1816 }
1817 }
1818 }
1819 }
1820 }
1821 }
1822 }
1823 }
1824 }
1825 }
1826 }
1827 }
1828 }
1829 }
1830 }
1831 }
1832 }
1833 }
1834 }
1835 }
1836 }
1837 }
1838 }
1839 }
1840 }
1841 }
1842 }
1843 }
1844 }
1845 }
1846 }
1847 }
1848 }
1849 }
1850 }
1851 }
1852 }
1853 }
1854 }
1855 }
1856 }
1857 }
1858 }
1859 }
1860 }
1861 }
1862 }
1863 }
1864 }
1865 }
1866 }
1867 }
1868 }
1869 }
1870 }
1871 }
1872 }
1873 }
1874 }
1875 }
1876 }
1877 }
1878 }
1879 }
1880 }
1881 }
1882 }
1883 }
1884 }
1885 }
1886 }
1887 }
1888 }
1889 }
1890 }
1891 }
1892 }
1893 }
1894 }
1895 }
1896 }
1897 }
1898 }
1899 }
1900 }
1901 }
1902 }
1903 }
1904 }
1905 }
1906 }
1907 }
1908 }
1909 }
1910 }
1911 }
1912 }
1913 }
1914 }
1915 }
1916 }
1917 }
1918 }
1919 }
1920 }
1921 }
1922 }
1923 }
1924 }
1925 }
1926 }
1927 }
1928 }
1929 }
1930 }
1931 }
1932 }
1933 }
1934 }
1935 }
1936 }
1937 }
1938 }
1939 }
1940 }
1941 }
1942 }
1943 }
1944 }
1945 }
1946 }
1947 }
1948 }
1949 }
1950 }
1951 }
1952 }
1953 }
1954 }
1955 }
1956 }
1957 }
1958 }
1959 }
1960 }
1961 }
1962 }
1963 }
1964 }
1965 }
1966 }
1967 }
1968 }
1969 }
1970 }
1971 }
1972 }
1973 }
1974 }
1975 }
1976 }
1977 }
1978 }
1979 }
1980 }
1981 }
1982 }
1983 }
1984 }
1985 }
1986 }
1987 }
1988 }
1989 }
1990 }
1991 }
1992 }
1993 }
1994 }
1995 }
1996 }
1997 }
1998 }
1999 }
2000 }

```

```

1494     ulong_t val;
1495     int xattr_allowed = dvp->v_vfsp->vfs_flag & VFS_XATTR;
1496     int sysattrs_allowed = 1;

1498     /*
1499      * We have to drop the lock on dvp. gfs_dir_create will
1500      * grab it for a VN_HOLD.
1501      */
1502     mutex_exit(&dvp->v_lock);

1504     /*
1505      * If dvp allows xattr creation, but not sysattr
1506      * creation, return the real xattr dir vp. We can't
1507      * use the vfs feature mask here because _PC_SATTR_ENABLED
1508      * has vnode-level granularity (e.g. .zfs).
1509      */
1510     error = VOP_PATHCONF(dvp, _PC_SATTR_ENABLED, &val, cr, NULL);
1511     if (error != 0 || val == 0)
1512         sysattrs_allowed = 0;

1514     if (!xattr_allowed && !sysattrs_allowed)
1515         return (EINVAL);

1517     if (!sysattrs_allowed) {
1518         struct pathname pn;
1519         char *nm = "";

1521         error = pn_get(nm, UIO_SYSSPACE, &pn);
1522         if (error)
1523             return (error);
1524         error = VOP_LOOKUP(dvp, nm, vpp, &pn,
1525             flags|LOOKUP_HAVE_SYSATTR_DIR, rootvp, cr, NULL,
1526             NULL, NULL);
1527         pn_free(&pn);
1528         return (error);
1529     }

1531     /*
1532      * Note that we act as if we were given CREATE_XATTR_DIR,
1533      * but only for creation of the GFS directory.
1534      */
1535     gfs_vp = gfs_dir_create(
1536         *vpp = gfs_dir_create(
1537             sizeof(xattr_dir_t), dvp, xattr_dir_ops, xattr_dirents,
1538             xattrdir_do_ino, MAXNAMELEN, NULL, xattr_lookup_cb);
1539     mutex_enter(&dvp->v_lock);
1540     if (dvp->v_xattrdir != NULL) {
1541         /*
1542          * We lost the race to create the xattr dir.
1543          * Destroy this one, use the winner. We can't
1544          * just call VN_RELE(*vpp), because the vnode
1545          * is only partially initialized.
1546          */
1547         gfs_dir_t *dp = gfs_vp->v_data;
1548         gfs_dir_t *dp = (*vpp)->v_data;

1550         ASSERT(gfs_vp->v_count == 1);
1551         vn_free(gfs_vp);
1552         ASSERT((*vpp)->v_count == 1);
1553         vn_free(*vpp);

1555         mutex_destroy(&dp->gfsd_lock);
1556         kmem_free(dp->gfsd_static,
1557             dp->gfsd_nstatic * sizeof(gfs_dirent_t));
1558         kmem_free(dp, dp->gfsd_file.gfs_size);

```

```

1553      /*
1554      * There is an implied VN_HOLD(dvp) here.  We should
1555      * be doing a VN_RELE(dvp) to clean up the reference
1556      * from gfs_vp, and then a VN_HOLD(dvp) for the new
1557      * from *vpp, and then a VN_HOLD(dvp) for the new
1558      * reference.  Instead, we just leave the count alone.
1559      */
1560      gfs_vp = dvp->v_xattrdir;
1561      VN_HOLD(gfs_vp);
1562      *vpp = dvp->v_xattrdir;
1563      VN_HOLD(*vpp);
1564      } else {
1565      gfs_vp->v_flag |= (V_XATTRDIR|V_SYSATTR);
1566      dvp->v_xattrdir = gfs_vp;
1567      (*vpp)->v_flag |= (V_XATTRDIR|V_SYSATTR);
1568      dvp->v_xattrdir = *vpp;
1569      }
1570      mutex_exit(&dvp->v_lock);
1571
1572      /*
1573      * In order to make this module relatively transparent
1574      * to the underlying filesystem, we need to lookup the
1575      * xattr dir in the lower filesystem and (if found)
1576      * keep a hold on it for as long as there is a hold
1577      * on the gfs_vp we're about to return.  This hold is
1578      * released in xattr_dir_inactive.
1579      */
1580      xattr_dir = gfs_vp->v_data;
1581      if ((dvp->v_vfsp->vfs_flag & VFS_XATTR) &&
1582          (xattr_dir->xattr_realvp == NULL)) {
1583          error = pn_get(nm, UIO_SYSSPACE, &pn);
1584          if (error == 0) {
1585              error = VOP_LOOKUP(dvp, nm, &real_vp, &pn,
1586                              flags|LOOKUP_HAVE_SYSATTR_DIR, rootvp, cr, NULL,
1587                              NULL, NULL);
1588              pn_free(&pn);
1589          }
1590          if (error == 0) {
1591              mutex_enter(&gfs_vp->v_lock);
1592              if (xattr_dir->xattr_realvp == NULL)
1593                  xattr_dir->xattr_realvp = real_vp;
1594              else
1595                  VN_RELE(real_vp);
1596              mutex_exit(&gfs_vp->v_lock);
1597          }
1598      }
1599      *vpp = gfs_vp;
1600      return (0);
1601      return (error);
1602  }
1603  }

```

unchanged portion omitted