

new/usr/src/cmd/zpool/zpool_main.c

```
*****
129759 Tue Apr 23 15:34:30 2013
new/usr/src/cmd/zpool/zpool_main.c
3745 zpool create should treat -o mountpoint and -m the same
Submitted by: Will Andrews <will@spectralogic.com>
Submitted by: Alan Somers <alans@spectralogic.com>
Reviewed by: Matthew Ahrens <mahrens@delphix.com>
*****
unchanged_portion_omitted
```

```
617 /*
618 * zpool create [-fnd] [-o property=value] ...
619 *           [-O file-system-property=value] ...
620 *           [-R root] [-m mountpoint] <pool> <dev> ...
621 *
622 *   -f      Force creation, even if devices appear in use
623 *   -n      Do not create the pool, but display the resulting layout if it
624 *          were to be created.
625 *   -R      Create a pool under an alternate root
626 *   -m      Set default mountpoint for the root dataset. By default it's
627 *          '/<pool>'
628 *   -o      Set property=value.
629 *   -d      Don't automatically enable all supported pool features
630 *          (individual features can be enabled with -o).
631 *   -O      Set fsproperty=value in the pool's root file system
632 *
633 * Creates the named pool according to the given vdev specification. The
634 * bulk of the vdev processing is done in get_vdev_spec() in zpool_vdev.c. Once
635 * we get the nvlist back from get_vdev_spec(), we either print out the contents
636 * (if '-n' was specified), or pass it to libzfs to do the creation.
637 */
638 int
639 zpool_do_create(int argc, char **argv)
640 {
641     boolean_t force = B_FALSE;
642     boolean_t dryrun = B_FALSE;
643     boolean_t enable_all_pool_feat = B_TRUE;
644     int c;
645     nvlist_t *nvroot = NULL;
646     char *poolname;
647     int ret = 1;
648     char *altroot = NULL;
649     char *mountpoint = NULL;
650     nvlist_t *fsprops = NULL;
651     nvlist_t *props = NULL;
652     char *propval;
653
654     /* check options */
655     while ((c = getopt(argc, argv, ":fndR:m:o:O:")) != -1) {
656         switch (c) {
657             case 'f':
658                 force = B_TRUE;
659                 break;
660             case 'n':
661                 dryrun = B_TRUE;
662                 break;
663             case 'd':
664                 enable_all_pool_feat = B_FALSE;
665                 break;
666             case 'R':
667                 altroot = optarg;
668                 if (add_prop_list(zpool_prop_to_name(
669                     ZPOOL_PROP_ALTROOT), optarg, &props, B_TRUE))
670                     goto errout;
671                 if (nvlist_lookup_string(props,
672                     zpool_prop_to_name(ZPOOL_PROP_CACHEFILE),
```

1

new/usr/src/cmd/zpool/zpool_main.c

```
673                                     &propval) == 0)
674             break;
675         if (add_prop_list(zpool_prop_to_name(
676             ZPOOL_PROP_CACHEFILE), "none", &props, B_TRUE))
677             goto errout;
678         break;
679     case 'm':
680         /* Equivalent to -O mountpoint=optarg */
681 #endif /* ! codereview */
682         mountpoint = optarg;
683         break;
684     case 'o':
685         if ((propval = strchr(optarg, '=')) == NULL) {
686             (void) fprintf(stderr, gettext("missing "
687                             "'=' for -o option\n"));
688             goto errout;
689         }
690         *propval = '\0';
691         propval++;
692         if (add_prop_list(optarg, propval, &props, B_TRUE))
693             goto errout;
694
695         /*
696          * If the user is creating a pool that doesn't support
697          * feature flags, don't enable any features.
698          */
699         if (zpool_name_to_prop(optarg) == ZPOOL_PROP_VERSION) {
700             char *end;
701             u_longlong_t ver;
702
703             ver = strtoull(propval, &end, 10);
704             if (*end == '\0' &&
705                 ver < SPA_VERSION_FEATURES) {
706                 enable_all_pool_feat = B_FALSE;
707             }
708             break;
709         }
710     case 'O':
711         if ((propval = strchr(optarg, '=')) == NULL) {
712             (void) fprintf(stderr, gettext("missing "
713                             "'=' for -O option\n"));
714             goto errout;
715         }
716         *propval = '\0';
717         propval++;
718
719         /*
720          * Mountpoints are checked and then added later.
721          * Uniquely among properties, they can be specified
722          * more than once, to avoid conflict with -m.
723          */
724         if (!strcmp(optarg,
725             zfs_prop_to_name(ZFS_PROP_MOUNTPOINT)))
726             mountpoint = propval;
727 #endif /* ! codereview */
728         if (add_prop_list(optarg, propval, &fsprops, B_FALSE))
729             goto errout;
730         break;
731     case ':':
732         (void) fprintf(stderr, gettext("missing argument for "
733                         "'%c' option\n"), optopt);
734         goto badusage;
735     case '?':
736         (void) fprintf(stderr, gettext("invalid option '%c'\n"),
737                         optopt);
738 }
```

2

```

739         goto badusage;
740     }
741
742     argc -= optind;
743     argv += optind;
744
745     /* get pool name and check number of arguments */
746     if (argc < 1) {
747         (void) fprintf(stderr, gettext("missing pool name argument\n"));
748         goto badusage;
749     }
750     if (argc < 2) {
751         (void) fprintf(stderr, gettext("missing vdev specification\n"));
752         goto badusage;
753     }
754
755     poolname = argv[0];
756
757     /*
758      * As a special case, check for use of '/' in the name, and direct the
759      * user to use 'zfs create' instead.
760      */
761     if (strchr(poolname, '/') != NULL) {
762         (void) fprintf(stderr, gettext("cannot create '%s': invalid "
763             "character '/' in pool name\n"), poolname);
764         (void) fprintf(stderr, gettext("use 'zfs create' to "
765             "create a dataset\n"));
766         goto errout;
767     }
768
769     /* pass off to get_vdev_spec for bulk processing */
770     nvroot = make_root_vdev(NULL, force, !force, B_FALSE, dryrun,
771     argc - 1, argv + 1);
772     if (nvroot == NULL)
773         goto errout;
774
775     /* make_root_vdev() allows 0 toplevel children if there are spares */
776     if (!zfs_allocatable_devs(nvroot)) {
777         (void) fprintf(stderr, gettext("invalid vdev "
778             "specification: at least one toplevel vdev must be "
779             "'specified'\n"));
780         goto errout;
781     }
782
783     if (altroot != NULL && altroot[0] != '/') {
784         (void) fprintf(stderr, gettext("invalid alternate root '%s': "
785             "'must be an absolute path\n"), altroot);
786         goto errout;
787     }
788
789     /*
790      * Check the validity of the mountpoint and direct the user to use the
791      * '-m' mountpoint option if it looks like its in use.
792      */
793     if (mountpoint == NULL ||
794         (strcmp(mountpoint, ZFS_MOUNTPOINT_LEGACY) != 0 &&
795         strcmp(mountpoint, ZFS_MOUNTPOINT_NONE) != 0)) {
796         char buf[MAXPATHLEN];
797         DIR *dirp;
798
799         if (mountpoint && mountpoint[0] != '/') {
800             (void) fprintf(stderr, gettext("invalid mountpoint "
801                 "'%s': must be an absolute path, 'legacy', or "
802                 "'none'\n"), mountpoint);
803             goto errout;
804         }

```

```

805     }
806
807     if (mountpoint == NULL) {
808         if (altroot != NULL)
809             (void) snprintf(buf, sizeof (buf), "%s/%s",
810                             altroot, poolname);
811         else
812             (void) snprintf(buf, sizeof (buf), "/%s",
813                             poolname);
814     } else {
815         if (altroot != NULL)
816             (void) snprintf(buf, sizeof (buf), "%s%s",
817                             altroot, mountpoint);
818         else
819             (void) snprintf(buf, sizeof (buf), "%s",
820                             mountpoint);
821     }
822
823     if ((dirp = opendir(buf)) == NULL && errno != ENOENT) {
824         (void) fprintf(stderr, gettext("mountpoint '%s' : "
825             "'%s'\n"), buf, strerror(errno));
826         (void) fprintf(stderr, gettext("use '-m' "
827             "option to provide a different default\n"));
828         goto errout;
829     } else if (dirp) {
830         int count = 0;
831
832         while (count < 3 && readdir(dirp) != NULL)
833             count++;
834         (void) closedir(dirp);
835
836         if (count > 2) {
837             (void) fprintf(stderr, gettext("mountpoint "
838                 "'%s' exists and is not empty\n"), buf);
839             (void) fprintf(stderr, gettext("use '-m' "
840                 "option to provide a "
841                 "different default\n"));
842             goto errout;
843     }
844
845     }
846
847     /*
848      * Now that the mountpoint's validity has been checked, ensure that
849      * the property is set appropriately prior to creating the pool.
850      */
851     if (mountpoint != NULL)
852         if (add_prop_list(zfs_prop_to_name(ZFS_PROP_MOUNTPOINT),
853                           mountpoint, &fsprops, B_FALSE))
854             goto errout;
855
856 #endif /* ! codereview */
857     if (dryrun) {
858         /*
859          * For a dry run invocation, print out a basic message and run
860          * through all the vdevs in the list and print out in an
861          * appropriate hierarchy.
862          */
863     }
864     (void) printf(gettext("would create '%s' with the "
865                     "following layout:\n\n"), poolname);
866
867     print_vdev_tree(NULL, poolname, nvroot, 0, B_FALSE);
868     if (num_logs(nvroot) > 0)
869         print_vdev_tree(NULL, "logs", nvroot, 0, B_TRUE);
870
871     ret = 0;

```

```
871         } else {
872             /*
873             * Hand off to libzfs.
874             */
875             if (enable_all_pool_feat) {
876                 int i;
877                 for (i = 0; i < SPA_FEATURES; i++) {
878                     char propname[MAXPATHLEN];
879                     zfeature_info_t *feat = &spa_feature_table[i];
880
881                     (void) snprintf(propname, sizeof (propname),
882                                     "feature@%s", feat->fi_uname);
883
884                     /*
885                     * Skip feature if user specified it manually
886                     * on the command line.
887                     */
888                     if (nvlist_exists(props, propname))
889                         continue;
890
891                     if (add_prop_list(propname, ZFS_FEATURE_ENABLED,
892                                     &props, B_TRUE) != 0)
893                         goto errout;
894
895             }
896             if (zpool_create(g_zfs, poolname,
897                             nvroot, props, fsprops) == 0) {
898                 zfs_handle_t *pool = zfs_open(g_zfs, poolname,
899                                              ZFS_TYPE_FILESYSTEM);
900                 if (pool != NULL) {
901                     if (mountpoint != NULL)
902                         verify(zfs_prop_set(pool,
903                                         zfs_prop_to_name(
904                                             ZFS_PROP_MOUNTPOINT),
905                                         mountpoint) == 0);
906                     if (zfs_mount(pool, NULL, 0) == 0)
907                         ret = zfs_shareall(pool);
908                     zfs_close(pool);
909                 }
910             } else if (libzfs_errno(g_zfs) == EZFS_INVALIDNAME) {
911                 (void) fprintf(stderr, gettext("pool name may have "
912                               "been omitted\n"));
913             }
914         }
915     }
916     errout:
917     nvlist_free(nvroot);
918     nvlist_free(fsprops);
919     nvlist_free(props);
920     return (ret);
921 }
```

unchanged portion omitted

new/usr/src/lib/libzfs/common/libzfs_pool.c

102992 Tue Apr 23 15:34:31 2013

new/usr/src/lib/libzfs/common/libzfs_pool.c

3745 zpool create should treat -O mountpoint and -m the same

Submitted by: Will Andrews <will@spectralogic.com>

Submitted by: Alan Somers <alans@spectralogic.com>

Reviewed by: Matthew Ahrens <mahrens@delphix.com>

_____ unchanged_portion_omitted _____

1070 /*
1071 * Create the named pool, using the provided vdev list. It is assumed
1072 * that the consumer has already validated the contents of the nvlist, so we
1073 * don't have to worry about error semantics.
1074 */
1075 int
1076 zpool_create(libzfs_handle_t *hdl, const char *pool, nvlist_t *nvroot,
1077 nvlist_t *props, nvlist_t *fsprops)
1078 {
1079 zfs_cmd_t zc = { 0 };
1080 nvlist_t *zc_fsprops = NULL;
1081 nvlist_t *zc_props = NULL;
1082 char msg[1024];
1083 char *altroot;
1084 int ret = -1;
1085
1086 (void) snprintf(msg, sizeof (msg), dgettext(TEXT_DOMAIN,
1087 "cannot create '%s'"), pool);
1088
1089 if (!zpool_name_valid(hdl, B_FALSE, pool))
1090 return (zfs_error(hdl, EZFS_INVALIDNAME, msg));
1091
1092 if (zcmd_write_conf_nvlist(hdl, &zc, nvroot) != 0)
1093 return (-1);
1094
1095 if (props) {
1096 prop_flags_t flags = { .create = B_TRUE, .import = B_FALSE };
1097
1098 if ((zc_props = zpool_valid_proplist(hdl, pool, props,
1099 SPA_VERSION_1, flags, msg)) == NULL)
1100 goto create_failed;
1101 }
1102
1103 if (fsprops) {
1104 uint64_t zoned;
1105 char *zonestr;
1106
1107 zoned = ((nvlist_lookup_string(fsprops,
1108 zfs_prop_to_name(ZFS_PROP_ZONEDE)), &zonestr) == 0) &&
1109 strcmp(zonestr, "on") == 0;
1110
1111 if ((zc_fsprops = zfs_valid_proplist(hdl,
1112 ZFS_TYPE_FILESYSTEM, fsprops, zoned, NULL, msg)) == NULL) {
1113 goto create_failed;
1114 }
1115 if (!zc_props &&
1116 (nvlist_alloc(&zc_props, NV_UNIQUE_NAME, 0) != 0)) {
1117 goto create_failed;
1118 }
1119 if (nvlist_add_nvlist(zc_props,
1120 ZPOOL_ROOTFS_PROPS, zc_fsprops) != 0) {
1121 goto create_failed;
1122 }
1123 }
1124 }

1

new/usr/src/lib/libzfs/common/libzfs_pool.c

1126 if (zc_props && zcmd_write_src_nvlist(hdl, &zc, zc_props) != 0)
1127 goto create_failed;
1128
1129 (void) strlcpy(zc.zc_name, pool, sizeof (zc.zc_name));
1130
1131 if ((ret = zfs_ioctl(hdl, ZFS_IOC_POOL_CREATE, &zc)) != 0) {
1132 zcmd_free_nvlists(&zc);
1133 nvlist_free(zc_props);
1134 nvlist_free(zc_fsprops);
1135
1136 switch (errno) {
1137 case EBUSY:
1138 /*
1139 * This can happen if the user has specified the same
1140 * device multiple times. We can't reliably detect this
1141 * until we try to add it and see we already have a
1142 * label.
1143 */
1144 zfs_error_aux(hdl, dgettext(TEXT_DOMAIN,
1145 "one or more vdevs refer to the same device"));
1146 return (zfs_error(hdl, EZFS_BADDEV, msg));
1147
1148 case EOVERRLOW:
1149 /*
1150 * This occurs when one of the devices is below
1151 * SPA_MINDEVSIZE. Unfortunately, we can't detect which
1152 * device was the problem device since there's no
1153 * reliable way to determine device size from userland.
1154 */
1155 {
1156 char buf[64];
1157
1158 zfs_nicenum(SPA_MINDEVSIZE, buf, sizeof (buf));
1159
1160 zfs_error_aux(hdl, dgettext(TEXT_DOMAIN,
1161 "one or more devices is less than the "
1162 "minimum size (%s)", buf));
1163 }
1164 return (zfs_error(hdl, EZFS_BADDEV, msg));
1165
1166 case ENOSPC:
1167 zfs_error_aux(hdl, dgettext(TEXT_DOMAIN,
1168 "one or more devices is out of space"));
1169 return (zfs_error(hdl, EZFS_BADDEV, msg));
1170
1171 case ENOTBLK:
1172 zfs_error_aux(hdl, dgettext(TEXT_DOMAIN,
1173 "cache device must be a disk or disk slice"));
1174 return (zfs_error(hdl, EZFS_BADDEV, msg));
1175
1176 default:
1177 return (zpool_standard_error(hdl, errno, msg));
1178 }
1179 }
1180
1181 /*
1182 * If this is an alternate root pool, then we automatically set the
1183 * mountpoint of the root dataset to be '/'.
1184 */
1185 if (nvlist_lookup_string(props, zpool_prop_to_name(ZPOOL_PROP_ALTROOT),
1186 &altroot) == 0) {
1187 zfs_handle_t *zhp;
1188
1189 verify((zhp = zfs_open(hdl, pool, ZFS_TYPE_DATASET)) != NULL);
1190 verify(zfs_prop_set(zhp, zfs_prop_to_name(ZFS_PROP_MOUNTPOINT),

2

```
1192             "/") == 0);  
1194         }  
1195     }  
  
1182     create_failed:  
1183         zcmd_free_nvlists(&zc);  
1184         nvlist_free(zc_props);  
1185         nvlist_free(zc_fsprops);  
1186         return (ret);  
1187 }  
unchanged portion omitted
```