

```

*****
60359 Wed Jan 30 11:17:45 2019
new/usr/src/cmd/zoneadmd/zoneadmd.c
10141 smatch fix for zoneadmd
*****
_____unchanged_portion_omitted_____

1827 int
1828 main(int argc, char *argv[])
1829 {
1830     int opt;
1831     zoneid_t zid;
1832     priv_set_t *privset;
1833     zone_state_t zstate;
1834     char parents_locale[MAXPATHLEN];
1835     brand_handle_t bh;
1836     int err;

1838     pid_t pid;
1839     sigset_t blockset;
1840     sigset_t block_cld;

1842     struct {
1843         sema_t sem;
1844         int status;
1845         zlog_t log;
1846     } *shstate;
1847     size_t shstatelen = getpagesize();

1849     zlog_t errlog;
1850     zlog_t *zlogp;

1852     int ctfd;

1854     progname = get_execbasename(argv[0]);

1856     /*
1857      * Make sure stderr is unbuffered
1858      */
1859     (void) setbuffer(stderr, NULL, 0);

1861     /*
1862      * Get out of the way of mounted filesystems, since we will daemonize
1863      * soon.
1864      */
1865     (void) chdir("/");

1867     /*
1868      * Use the default system umask per PSARC 1998/110 rather than
1869      * anything that may have been set by the caller.
1870      */
1871     (void) umask(CMASK);

1873     /*
1874      * Initially we want to use our parent's locale.
1875      */
1876     (void) setlocale(LC_ALL, "");
1877     (void) textdomain(TEXT_DOMAIN);
1878     (void) strlcpy(parents_locale, setlocale(LC_MESSAGES, NULL),
1879                 sizeof (parents_locale));

1881     /*
1882      * This zlog_t is used for writing to stderr
1883      */
1884     errlog.logfile = stderr;
1885     errlog.buflen = errlog.loglen = 0;

```

```

1886     errlog.buf = errlog.log = NULL;
1887     errlog.locale = parents_locale;

1889     /*
1890      * We start off writing to stderr until we're ready to daemonize.
1891      */
1892     zlogp = &errlog;

1894     /*
1895      * Process options.
1896      */
1897     while ((opt = getopt(argc, argv, "R:z:")) != EOF) {
1898         switch (opt) {
1899             case 'R':
1900                 zonecfg_set_root(optarg);
1901                 break;
1902             case 'z':
1903                 zone_name = optarg;
1904                 break;
1905             default:
1906                 usage();
1907         }
1908     }

1910     if (zone_name == NULL)
1911         usage();

1913     /*
1914      * Because usage() prints directly to stderr, it has gettext()
1915      * wrapping, which depends on the locale. But since zerror() calls
1916      * localize() which tweaks the locale, it is not safe to call zerror()
1917      * until after the last call to usage(). Fortunately, the last call
1918      * to usage() is just above and the first call to zerror() is just
1919      * below. Don't mess this up.
1920      */
1921     if (strcmp(zone_name, GLOBAL_ZONENAME) == 0) {
1922         zerror(zlogp, B_FALSE, "cannot manage the %s zone",
1923             GLOBAL_ZONENAME);
1924         return (1);
1925     }

1927     if (zone_get_id(zone_name, &zid) != 0) {
1928         zerror(zlogp, B_FALSE, "could not manage %s: %s", zone_name,
1929             zonecfg_strerror(Z_NO_ZONE));
1930         return (1);
1931     }

1933     if ((err = zone_get_state(zone_name, &zstate)) != Z_OK) {
1934         zerror(zlogp, B_FALSE, "failed to get zone state: %s",
1935             zonecfg_strerror(err));
1936         return (1);
1937     }
1938     if (zstate < ZONE_STATE_INCOMPLETE) {
1939         zerror(zlogp, B_FALSE,
1940             "cannot manage a zone which is in state '%s'",
1941             zone_state_str(zstate));
1942         return (1);
1943     }

1945     if (zonecfg_default_brand(default_brand,
1946         sizeof (default_brand)) != Z_OK) {
1947         zerror(zlogp, B_FALSE, "unable to determine default brand");
1948         return (1);
1949     }

1951     /* Get a handle to the brand info for this zone */

```

```

1952     if (zone_get_brand(zone_name, brand_name, sizeof (brand_name))
1953         != Z_OK) {
1954         zerror(zlogp, B_FALSE, "unable to determine zone brand");
1955         return (1);
1956     }
1957     zone_isnative = (strcmp(brand_name, NATIVE_BRAND_NAME) == 0);
1958     zone_islabeled = (strcmp(brand_name, LABELED_BRAND_NAME) == 0);

1960     /*
1961     * In the alternate root environment, the only supported
1962     * operations are mount and unmount. In this case, just treat
1963     * the zone as native if it is cluster. Cluster zones can be
1964     * native for the purpose of LU or upgrade, and the cluster
1965     * brand may not exist in the miniroot (such as in net install
1966     * upgrade).
1967     */
1968     if (strcmp(brand_name, CLUSTER_BRAND_NAME) == 0) {
1969         zone_iscluster = B_TRUE;
1970         if (zonecfg_in_alt_root()) {
1971             (void) strcpy(brand_name, default_brand,
1972                 sizeof (brand_name));
1973         }
1974     } else {
1975         zone_iscluster = B_FALSE;
1976     }

1978     if ((bh = brand_open(brand_name)) == NULL) {
1979         zerror(zlogp, B_FALSE, "unable to open zone brand");
1980         return (1);
1981     }

1983     /* Get state change brand hooks. */
1984     if (brand_callback_init(bh, zone_name) == -1) {
1985         zerror(zlogp, B_TRUE,
1986             "failed to initialize brand state change hooks");
1987         brand_close(bh);
1988         return (1);
1989     }

1991     brand_close(bh);

1993     /*
1994     * Check that we have all privileges. It would be nice to pare
1995     * this down, but this is at least a first cut.
1996     */
1997     if ((privset = priv_allocset()) == NULL) {
1998         zerror(zlogp, B_TRUE, "%s failed", "priv_allocset");
1999         return (1);
2000     }

2002     if (getppriv(PRIV_EFFECTIVE, privset) != 0) {
2003         zerror(zlogp, B_TRUE, "%s failed", "getppriv");
2004         priv_freerset(privset);
2005         return (1);
2006     }

2008     if (priv_isfullset(privset) == B_FALSE) {
2009         zerror(zlogp, B_FALSE, "You lack sufficient privilege to "
2010             "run this command (all privs required)");
2011         priv_freerset(privset);
2012         return (1);
2013     }
2014     priv_freerset(privset);

2016     if (mkzonedir(zlogp) != 0)
2017         return (1);

```

```

2019     /*
2020     * Pre-fork: setup shared state
2021     */
2022     if ((shstate = (void *)mmap(NULL, shstatelen,
2023         PROT_READ|PROT_WRITE, MAP_SHARED|MAP_ANON, -1, (off_t)0)) ==
2024         MAP_FAILED) {
2025         zerror(zlogp, B_TRUE, "%s failed", "mmap");
2026         return (1);
2027     }
2028     if (sema_init(&shstate->sem, 0, USYNC_PROCESS, NULL) != 0) {
2029         zerror(zlogp, B_TRUE, "%s failed", "sema_init()");
2030         (void) munmap((char *)shstate, shstatelen);
2031         return (1);
2032     }
2033     shstate->log.logfile = NULL;
2034     shstate->log.bufalen = shstatelen - sizeof (*shstate);
2035     shstate->log.loglen = shstate->log.bufalen;
2036     shstate->log.buf = (char *)shstate + sizeof (*shstate);
2037     shstate->log.log = shstate->log.buf;
2038     shstate->log.locale = parents_locale;
2039     shstate->status = -1;

2041     /*
2042     * We need a SIGCHLD handler so the sema_wait() below will wake
2043     * up if the child dies without doing a sema_post().
2044     */
2045     (void) sigset(SIGCHLD, sigchld);
2046     /*
2047     * We must mask SIGCHLD until after we've coped with the fork
2048     * sufficiently to deal with it; otherwise we can race and
2049     * receive the signal before pid has been initialized
2050     * (yes, this really happens).
2051     */
2052     (void) sigemptyset(&block_cld);
2053     (void) sigaddset(&block_cld, SIGCHLD);
2054     (void) sigprocmask(SIG_BLOCK, &block_cld, NULL);

2056     /*
2057     * The parent only needs stderr after the fork, so close other fd's
2058     * that we inherited from zoneadm so that the parent doesn't have those
2059     * open while waiting. The child will close the rest after the fork.
2060     */
2061     closefrom(3);

2063     if ((ctfd = init_template()) == -1) {
2064         zerror(zlogp, B_TRUE, "failed to create contract");
2065         return (1);
2066     }

2068     /*
2069     * Do not let another thread localize a message while we are forking.
2070     */
2071     (void) mutex_lock(&msglock);
2072     pid = fork();
2073     (void) mutex_unlock(&msglock);

2075     /*
2076     * In all cases (parent, child, and in the event of an error) we
2077     * don't want to cause creation of contracts on subsequent fork()s.
2078     */
2079     (void) ct_tmpl_clear(ctfd);
2080     (void) close(ctfd);

2082     if (pid == -1) {
2083         zerror(zlogp, B_TRUE, "could not fork");

```

```

2084         return (1);
2086     } else if (pid > 0) { /* parent */
2087         (void) sigprocmask(SIG_UNBLOCK, &block_cld, NULL);
2088         /*
2089          * This marks a window of vulnerability in which we receive
2090          * the SIGCLD before falling into sema_wait (normally we would
2091          * get woken up from sema_wait with EINTR upon receipt of
2092          * SIGCLD). So we may need to use some other scheme like
2093          * sema_posting in the sigcld handler.
2094          * blech
2095          */
2096         (void) sema_wait(&shstate->sem);
2097         (void) sema_destroy(&shstate->sem);
2098         if (shstate->status != 0)
2099             (void) waitpid(pid, NULL, WNOHANG);
2100         /*
2101          * It's ok if we die with SIGPIPE. It's not like we could have
2102          * done anything about it.
2103          */
2104         (void) fprintf(stderr, "%s", shstate->log.buf);
2105         _exit(shstate->status == 0 ? 0 : 1);
2106     }
2108     /*
2109     * The child charges on.
2110     */
2111     (void) sigset(SIGCHLD, SIG_DFL);
2112     (void) sigprocmask(SIG_UNBLOCK, &block_cld, NULL);
2114     /*
2115     * SIGPIPE can be delivered if we write to a socket for which the
2116     * peer endpoint is gone. That can lead to too-early termination
2117     * of zoneadmd, and that's not good eats.
2118     */
2119     (void) sigset(SIGPIPE, SIG_IGN);
2120     /*
2121     * Stop using stderr
2122     */
2123     zlogp = &shstate->log;
2125     /*
2126     * We don't need stdout/stderr from now on.
2127     */
2128     closefrom(0);
2130     /*
2131     * Initialize the syslog zlog_t. This needs to be done after
2132     * the call to closefrom().
2133     */
2134     logsys.buf = logsys.log = NULL;
2135     logsys buflen = logsys.loglen = 0;
2136     logsys.logfile = NULL;
2137     logsys.locale = DEFAULT_LOCALE;
2139     openlog("zoneadmd", LOG_PID, LOG_DAEMON);
2141     /*
2142     * The eventstream is used to publish state changes in the zone
2143     * from the door threads to the console I/O poller.
2144     */
2145     if (eventstream_init() == -1) {
2146         zerror(zlogp, B_TRUE, "unable to create eventstream");
2147         goto child_out;
2148     }

```

```

2150         (void) snprintf(zone_door_path, sizeof (zone_door_path),
2151             "%s" ZONE_DOOR_PATH, zonecfg_get_root(), zone_name);
2153     /*
2154     * See if another zoneadmd is running for this zone. If not, then we
2155     * can now modify system state.
2156     */
2157     if (make_daemon_exclusive(zlogp) == -1)
2158         goto child_out;
2161     /*
2162     * Create/join a new session; we need to be careful of what we do with
2163     * the console from now on so we don't end up being the session leader
2164     * for the terminal we're going to be handing out.
2165     */
2166     (void) setsid();
2168     /*
2169     * This thread shouldn't be receiving any signals; in particular,
2170     * SIGCHLD should be received by the thread doing the fork().
2171     */
2172     (void) sigfillset(&blockset);
2173     (void) thr_sigsetmask(SIG_BLOCK, &blockset, NULL);
2175     /*
2176     * Setup the console device and get ready to serve the console;
2177     * once this has completed, we're ready to let console clients
2178     * make an attempt to connect (they will block until
2179     * serve_console_sock() below gets called, and any pending
2180     * connection is accept()ed).
2181     */
2182     if (!zonecfg_in_alt_root() && init_console(zlogp) < 0)
2183         goto child_out;
2185     /*
2186     * Take the lock now, so that when the door server gets going, we
2187     * are guaranteed that it won't take a request until we are sure
2188     * that everything is completely set up. See the child_out: label
2189     * below to see why this matters.
2190     */
2191     (void) mutex_lock(&lock);
2193     /* Init semaphore for scratch zones. */
2194     if (sema_init(&scratch_sem, 0, USYNC_THREAD, NULL) == -1) {
2195         zerror(zlogp, B_TRUE,
2196             "failed to initialize semaphore for scratch zone");
2197         goto child_out;
2198     }
2200     /* open the dladm handle */
2201     if (dladm_open(&dld_handle) != DLADM_STATUS_OK) {
2202         zerror(zlogp, B_FALSE, "failed to open dladm handle");
2203         goto child_out;
2204     }
2206     /*
2207     * Note: door setup must occur *after* the console is setup.
2208     * This is so that as zlogin tests the door to see if zoneadmd
2209     * is ready yet, we know that the console will get serviced
2210     * once door_info() indicates that the door is "up".
2211     */
2212     if (setup_door(zlogp) == -1)
2213         goto child_out;
2215     /*

```

```

2216     * Things seem OK so far; tell the parent process that we're done
2217     * with setup tasks. This will cause the parent to exit, signalling
2218     * to zoneadm, zlogin, or whatever forked it that we are ready to
2219     * service requests.
2220     */
2221     shstate->status = 0;
2222     (void) sema_post(&shstate->sem);
2223     (void) munmap((char *)shstate, shstatelen);
2224     shstate = NULL;

2226     (void) mutex_unlock(&lock);

2228     /*
2229     * zlogp is now invalid, so reset it to the syslog logger.
2230     */
2231     zlogp = &logsys;

2233     /*
2234     * Now that we are free of any parents, switch to the default locale.
2235     */
2236     (void) setlocale(LC_ALL, DEFAULT_LOCALE);

2238     /*
2239     * At this point the setup portion of main() is basically done, so
2240     * we reuse this thread to manage the zone console. When
2241     * serve_console() has returned, we are past the point of no return
2242     * in the life of this zoneadmd.
2243     */
2244     if (zonecfg_in_alt_root()) {
2245         /*
2246         * This is just awful, but mounted scratch zones don't (and
2247         * can't) have consoles. We just wait for unmount instead.
2248         */
2249         while (sema_wait(&scratch_sem) == EINTR)
2250             ;
2251     } else {
2252         serve_console(zlogp);
2253         assert(in_death_throes);
2254     }

2256     /*
2257     * This is the next-to-last part of the exit interlock. Upon calling
2258     * fdetach(), the door will go unreferenced; once any
2259     * outstanding requests (like the door thread doing Z_HALT) are
2260     * done, the door will get an UNREF notification; when it handles
2261     * the UNREF, the door server will cause the exit. It's possible
2262     * that fdetach() can fail because the file is in use, in which
2263     * case we'll retry the operation.
2264     */
2265     assert(!MUTEX_HELD(&lock));
2266     for (;;) {
2267         if ((fdetach(zone_door_path) == 0) || (errno != EBUSY))
2268             break;
2269         yield();
2270     }

2272     for (;;)
2273         (void) pause();

2275 child_out:
2276     assert(pid == 0);

2277     if (shstate != NULL) {
2278         shstate->status = -1;
2279         (void) sema_post(&shstate->sem);
2280         (void) munmap((char *)shstate, shstatelen);

```

```

2281     }

2282     /*
2283     * This might trigger an unref notification, but if so,
2284     * we are still holding the lock, so our call to exit will
2285     * ultimately win the race and will publish the right exit
2286     * code.
2287     */
2288     if (zone_door != -1) {
2289         assert(MUTEX_HELD(&lock));
2290         (void) door_revoke(zone_door);
2291         (void) fdetach(zone_door_path);
2292     }

2294     if (dld_handle != NULL)
2295         dladm_close(dld_handle);

2297     return (1); /* return from main() forcibly exits an MT process */
2298 }

```

unchanged portion omitted