

```
*****
9672 Fri Apr 1 09:11:12 2016
new/usr/src/lib/libdns_sd/common/dnssd_clientstub.c
6771 end-of-loop code not reached in common/dnssd_clientstub.c
*****  

_____ unchanged_portion_omitted _____  

596 #define deliver_request_bailout(MSG) \
597     syslog(LOG_WARNING, "dnssd_clientstub deliver_request: %s failed %d (%s)", \
597     do { syslog(LOG_WARNING, "dnssd_clientstub deliver_request: %s failed %d (%s  

598 static DNSServiceErrorType deliver_request(ipc_msg_hdr *hdr, DNSServiceOp *sdr)  

599 {  

600     uint32_t datalen = hdr->datalen; // We take a copy here because we're going  

601     #if defined(USE_TCP_LOOPBACK) || defined(USE_NAMED_ERROR_RETURN_SOCKET)  

602     char *const data = (char *)hdr + sizeof(ipc_msg_hdr);  

603     #endif  

604     dnssd_sock_t listenfd = dnssd_InvalidSocket, errsd = dnssd_InvalidSocket;  

605     DNSServiceErrorType err = kDNSServiceErr_Unknown; // Default for the "goto  

606     int MakeSeparateReturnSocket = 0;  

607  

608     // Note: need to check hdr->op, not sdr->op.  

609     // hdr->op contains the code for the specific operation we're currently doing  

610     // contains the original parent DNSServiceOp (e.g. for an add_record_request  

611     // add_record_request but the parent sdr->op will be connection_request or r  

612     if (sdr->primary ||  

613         hdr->op == reg_record_request || hdr->op == add_record_request || hdr->op  

614         MakeSeparateReturnSocket = 1;  

615  

616     if (!DNSServiceRefValid(sdr))  

617     {  

618         if (hdr)  

619             free(hdr);  

620         syslog(LOG_WARNING, "dnssd_clientstub deliver_request: invalid DNSServiceOp");  

621         return kDNSServiceErr_BadReference;  

622     }  

623  

624     if (!hdr)  

625     {  

626         syslog(LOG_WARNING, "dnssd_clientstub deliver_request: !hdr");  

627         return kDNSServiceErr_Unknown;  

628     }  

629  

630     if (MakeSeparateReturnSocket)  

631     {  

632         #if defined(USE_TCP_LOOPBACK)  

633         {  

634             union { uint16_t s; u_char b[2]; } port;  

635             dnssd_sockaddr_t caddr;  

636             dnssd_socklen_t len = (dnssd_socklen_t) sizeof(caddr);  

637             listenfd = socket(AF_DNSSD, SOCK_STREAM, 0);  

638             if (!dnssd_SocketValid(listenfd)) {  

639                 deliver_request_bailout("TCP socket");  

640             }  

641             if (!dnssd_SocketValid(listenfd)) deliver_request_bailout("TCP socket");  

642  

643             caddr.sin_family = AF_INET;  

644             caddr.sin_port = 0;  

645             caddr.sin_addr.s_addr = inet_addr(MDNS_TCP_SERVERADDR);  

646             if (bind(listenfd, (struct sockaddr*) &caddr, sizeof(caddr)) < 0) {  

647                 deliver_request_bailout("TCP bind");  

648             }  

649             if (getsockname(listenfd, (struct sockaddr*) &caddr, &len) < 0) {  

650                 deliver_request_bailout("TCP getsockname");  

651             }  

652             if (listen(listenfd, 1)
```

```
        deliver_request_bailout("TCP listen");  

        }  

        if (bind(listenfd, (struct sockaddr*) &caddr, sizeof(caddr)) < 0) {  

            if (getsockname(listenfd, (struct sockaddr*) &caddr, &len) < 0) {  

                if (listen(listenfd, 1) < 0) {  

                    port.s = caddr.sin_port;  

                    data[0] = port.b[0]; // don't switch the byte order, as the  

                    data[1] = port.b[1]; // daemon expects it in network byte order  

                }  

            }#elif defined(USE_NAMED_ERROR_RETURN_SOCKET)  

            {  

                mode_t mask;  

                int bindresult;  

                dnssd_sockaddr_t caddr;  

                listenfd = socket(AF_DNSSD, SOCK_STREAM, 0);  

                if (!dnssd_SocketValid(listenfd)) {  

                    deliver_request_bailout("USE_NAMED_ERROR_RETURN_SOCKET");  

                }  

                if (!dnssd_SocketValid(listenfd)) deliver_request_bailout("USE_NAMED_ERROR_RETURN_SOCKET");  

                caddr.sun_family = AF_LOCAL;  

                // According to Stevens (section 3.2), there is no portable way to  

                // determine whether sa_len is defined on a particular platform.  

                #ifndef NOT_HAVE_SA_LEN  

                caddr.sun_len = sizeof(struct sockaddr_un);  

                #endif  

                strcpy(caddr.sun_path, data);  

                mask = umask(0);  

                bindresult = bind(listenfd, (struct sockaddr*) &caddr, sizeof(caddr));  

                umask(mask);  

                if (bindresult < 0) {  

                    deliver_request_bailout("USE_NAMED_ERROR_RETURN_SOCKET");  

                }  

                if (bindresult < 0) deliver_request_bailout("USE_NAMED_ERROR_RETURN_SOCKET");  

                if (listen(listenfd, 1) < 0) {  

                    deliver_request_bailout("USE_NAMED_ERROR_RETURN_SOCKET");  

                }  

            }#else  

            {  

                dnssd_sock_t sp[2];  

                if (socketpair(AF_DNSSD, SOCK_STREAM, 0, sp) < 0) {  

                    deliver_request_bailout("socketpair");  

                }  

                if (socketpair(AF_DNSSD, SOCK_STREAM, 0, sp) < 0) deliver_request_bailout("socketpair");  

                else {  

                    errsd = sp[0]; // We'll read our four-byte error code from  

                    listenfd = sp[1]; // We'll send sp[1] to the daemon  

                    #if !defined(__ppc__) & defined(SO_DEFUNCTOK)  

                    {  

                        int defunct = 1;  

                        if (setsockopt(errsd, SOL_SOCKET, SO_DEFUNCTOK, &defunct, si)  

                            syslog(LOG_WARNING, "dnssd_clientstub ConnectToServer: socketpair error");  

                    }#endif  

                }  

            }#endif  

        }#endif  

    }#if defined(USE_TCP_LOOPBACK) && !defined(USE_NAMED_ERROR_RETURN_SOCKET)  

    // If we're going to make a separate error return socket, and pass it to the  

    // using sendmsg, then we'll hold back one data byte to go with it.  

    // On some versions of Unix (including Leopard) sending a control message will
```

```

712     // any associated data does not work reliably -- e.g. one particular issue w
713     // into is that if the receiving program is in a kqueue loop waiting to be n
714     // of the received message, it doesn't get woken up when the control message
715     if (MakeSeparateReturnSocket || sdr->op == send_bpf)
716         datalen--; // Okay to use sdr->op when checking for op == send_bpf
717 #endiff

719     // At this point, our listening socket is set up and waiting, if necessary,
720     ConvertHeaderBytes(hdr);
721     //syslog(LOG_WARNING, "dnssd_clientstub deliver_request writing %lu bytes",
722     //if (MakeSeparateReturnSocket) syslog(LOG_WARNING, "dnssd_clientstub deliver
723 #if TEST_SENDING_ONE_BYTE_AT_A_TIME
724     unsigned int i;
725     for (i=0; i<datalen + sizeof(ipc_msg_hdr); i++)
726     {
727         syslog(LOG_WARNING, "dnssd_clientstub deliver_request writing %d", i);
728         if (write_all(sdr->sockfd, ((char *)hdr)+i, 1) < 0)
729             { syslog(LOG_WARNING, "write_all (byte %u) failed", i); goto cleanup; }
730         usleep(10000);
731     }
732 #else
733     if (write_all(sdr->sockfd, (char *)hdr, datalen + sizeof(ipc_msg_hdr)) < 0)
734     {
735         // write_all already prints an error message if there is an error writin
736         // the socket except for DEFUNCT. Logging here is unnecessary and also w
737         // in the case of DEFUNCT sockets
738         syslog(LOG_INFO, "dnssd_clientstub deliver_request ERROR: write_all(%d,
739                 sdr->sockfd, (unsigned long)(datalen + sizeof(ipc_msg_hdr)));
740         goto cleanup;
741     }
742 #endiff

744     if (!MakeSeparateReturnSocket)
745         errsd = sdr->sockfd;
746     if (MakeSeparateReturnSocket || sdr->op == send_bpf) // Okay to use sdr->
747     {
748 #if defined(USE_TCP_LOOPBACK) || defined(USE_NAMED_ERROR_RETURN_SOCKET)
749         // At this point we may wait in accept for a few milliseconds waiting fo
750         // but that's okay -- the daemon should not take more than a few millise
751         // set_waitlimit() ensures we do not block indefinitely just in case som
752         dnssd_sockaddr_t daddr;
753         dnssd_socklen_t len = sizeof(daddr);
754         if ((err = set_waitlimit(listenfd, DNSSD_CLIENT_TIMEOUT)) != kDNSService
755             goto cleanup;
756         errsd = accept(listenfd, (struct sockaddr *)&daddr, &len);
757         if (!dnssd_SocketValid(errsd))
758             if (!dnssd_SocketValid(errsd))
759                 deliver_request_bailout("accept");
760 #else

762         struct iovec vec = { ((char *)hdr) + sizeof(ipc_msg_hdr) + datalen, 1 };
763         struct msghdr msg;
764         struct cmsghdr *cmsg;
765         char cbuf[CMSG_SPACE(4 * sizeof(dnssd_sock_t))];

767         msg.msg_name      = 0;
768         msg.msg_namelen   = 0;
769         msg.msg_iov        = &vec;
770         msg.msg_iolen     = 1;
771         msg.msg_flags      = 0;
772         if (MakeSeparateReturnSocket || sdr->op == send_bpf) // Okay to use s
773         {
774             if (sdr->op == send_bpf)
775             {
776                 int i;

```

```

777         char p[12]; // Room for "/dev/bpf999" with terminating null
778         for (i=0; i<100; i++)
779         {
780             snprintf(p, sizeof(p), "/dev/bpf%d", i);
781             listenfd = open(p, O_RDWR, 0);
782             //if (!dnssd_SocketValid(listenfd)) syslog(LOG_WARNING, "Send
783             if (!dnssd_SocketValid(listenfd) && dnssd_errno != EBUSY)
784                 syslog(LOG_WARNING, "Error opening %s %d (%s)", p, dnssd
785                 if (dnssd_SocketValid(listenfd) || dnssd_errno != EBUSY) bre
786             }
787             msg.msg_control    = cbuf;
788             msg.msg_controllen = CMSG_LEN(sizeof(dnssd_sock_t));
789
790             cmsg = CMSG_FIRSTHDR(&msg);
791             cmsg->cmsg_len    = CMSG_LEN(sizeof(dnssd_sock_t));
792             cmsg->cmsg_level  = SOL_SOCKET;
793             cmsg->cmsg_type   = SCM_RIGHTS;
794             *((dnssd_sock_t *)CMSG_DATA(cmsg)) = listenfd;
795         }
796
797 #if TEST_KQUEUE_CONTROL_MESSAGE_BUG
798         sleep(1);
799 #endiff

800 #if DEBUG_64BIT_SCm_RIGHTS
801     syslog(LOG_WARNING, "dnssd_clientstub sendmsg read sd=%d write sd=%d %ld
802             errsd, listenfd, sizeof(dnssd_sock_t), sizeof(void*),
803             sizeof(struct cmsghdr) + sizeof(dnssd_sock_t),
804             CMSG_LEN(sizeof(dnssd_sock_t)), (long)CMSG_SPACE(sizeof(dnssd_soc
805             (long)((char*)CMSG_DATA(cmsg) + 4 - cbuf));
806 #endiff // DEBUG_64BIT_SCm_RIGHTS

807         if (sendmsg(sdr->sockfd, &msg, 0) < 0)
808         {
809             syslog(LOG_WARNING, "dnssd_clientstub deliver_request ERROR: sendmsg
810                 errsd, listenfd, dnssd_errno, dnssd_strerror(dnssd_errno));
811             err = kDNSServiceErr_Incompatible;
812             goto cleanup;
813         }
814
815 #if DEBUG_64BIT_SCm_RIGHTS
816         syslog(LOG_WARNING, "dnssd_clientstub sendmsg read sd=%d write sd=%d oka
817 #endiff // DEBUG_64BIT_SCm_RIGHTS

818 #endiff
819         // Close our end of the socketpair *before* calling read_all() to get th
820         // Otherwise, if the daemon closes our socket (or crashes), we will have
821         // in read_all() because the socket is not closed (we still have an open
822         // Note: listenfd is overwritten in the case of send_bpf above and that
823         // for send_bpf operation.
824         dnssd_close(listenfd);
825         listenfd = dnssd_InvalidSocket; // Make sure we don't close it a second
826
827
828
829
830 }

831
832         // At this point we may wait in read_all for a few milliseconds waiting for
833         // but that's okay -- the daemon should not take more than a few millisecond
834         // set_waitlimit() ensures we do not block indefinitely just in case somethi
835         if (sdr->op == send_bpf) // Okay to use sdr->op when checking for op == s
836             err = kDNSServiceErr_NoError;
837         else if ((err = set_waitlimit(errsd, DNSSD_CLIENT_TIMEOUT)) == kDNSServiceEr
838         {
839             if (read_all(errsd, (char*)&err, (int)sizeof(err)) < 0)
840                 err = kDNSServiceErr_ServiceNotRunning; // On failure read_all will
841             else
842                 err = ntohl(err);

```

```
843     }
844     //syslog(LOG_WARNING, "dnssd_clientstub deliver_request: retrieved error cod
846 cleanup:
847     if (MakeSeparateReturnSocket)
848     {
849         if (dnssd_SocketValid(listenfd)) dnssd_close(listenfd);
850         if (dnssd_SocketValid(errsd)) dnssd_close(errsd);
851 #if defined(USE_NAMED_ERROR_RETURN_SOCKET)
852         // syslog(LOG_WARNING, "dnssd_clientstub deliver_request: removing UDS:
853         if (unlink(data) != 0)
854             syslog(LOG_WARNING, "dnssd_clientstub WARNING: unlink(\"%s\") failed
855         // else syslog(LOG_WARNING, "dnssd_clientstub deliver_request: removed U
856 #endif
857     }
858     free(hdr);
859     return err;
860 }
unchanged portion omitted
```