

Interface-Related Changes to libresolv2

How these were implemented in the port to ISC libbind(9)6.0

For **Sun Microsystems**

By **Ed Posnak, Venev Inc.**

415.254.0086

www.venev.com

version 1.0

July 1, 2009

Interface-Related Changes

The following list enumerates all the changes that were incorporated into ISC code for the libbind(9)6.0 port.

Interface Changes

1. Some function names conflict with libnsl. Makefile.com renames these as:

```
gethostbyname=res_gethostbyname
gethostbyaddr=res_gethostbyaddr
getnetbyname=res_getnetbyname
gethostbyname2=res_gethostbyname2
getnetbyaddr=res_getnetbyaddr
sethostent=res_sethostent
endhostent=res_endhostent
gethostent=res_gethostent
setnetent=res_setnetent
endnetent=res_endnetent
getnetent=res_getnetent
socket=_socket
getipnodebyname=res_getipnodebyname
getipnodebyaddr=res_getipnodebyaddr
freehostent=res_freehostent
getaddrinfo=res_getaddrinfo
freeaddrinfo=res_freeaddrinfo
```

NOTE: many of these also have different signatures from OpenSolaris and have workarounds in the ISC code. We can use the unified adapter strategy (see GeneralPortingStrategy.doc) to get rid of all of those changes.

Category: Solaris-Specific

Action: renamed all funcs in port_netdb.h, Left socket in Makefile.com, i.e. SOLCOMPAT=-Dsocket=_socket.

Files:

```
Makefile.com
```

2. Modified llib-lresolv to more closely reflect the actual resolver interface.

NOTE: llib-resolv reflects only a subset of the functions defined in the public headers.

Category: Solaris-Specific

Action: left llib-lresolv alone; it's not an ISC file.

Files:

```
llib-lresolv
```

3. For BIND 8.2, ISC changed the ns_updrec struct, which, in Solaris, is defined in <arpa/nameser.h> . To retain the original interface, the ISC functions (res_update, res_mkupdate, etc.) were renamed in res_update.h and wrapped by functions defined in sunw_updrec.c.

Category: Solaris-Specific

Action: updated sunw_updrec.c. Left res_nmkupdate and res_nupdate alone.

Files:

```
sunw/sunw_updrec.c
include/res_update.h
```

4. Dummy function (never called) to avoid changing the mapfile depending on whether or not WANT_IRS_NIS is set.

Category: Solaris-Specific

Action: dropped. irs is now part of libresolv2

Files:

```
sunw/sunw_irs_nis_acc.c
irs/Makefile.com
```

5. Some function return 'int' on Solaris, but 'void' in ISC libbind.

```
setnetgrent(3C)/endnetgrent(3C)
setnetent(3SOCKET)/endnetent(3SOCKET)
setprotoent(3SOCKET)/endprotoent(3SOCKET)
setservent(3C)/endservent(3C)
sethostent(3NSL) and endhostent(3NSL)
```

NOTE: Most of these are already being renamed to res_* versions to avoid conflict with libnsl.

Category: Solaris-Specific

Action: renamed all functions to res_* in headers outside of ISC code.

Files:

```
irs/getnetgrent.c
irs/getnetent.c
irs/getprotoent.c
irs/getservent.c
irs/gethostent.c
include/port_after.h
```

6. SUNW_LIBNSL: Used to remove some routines that are implemented by the Solaris libnsl.so.1; see 4197284.

Category: Solaris-Specific

Action: removed .o targets from INET_OBJS in Makefile.com. SUNW_LIBNSL no longer needed in ISC code.

Files:

```
inet/inet_ntoa.c
inet/inet_ntop.c
inet/inet_pton.c
```

7. Some 'pragma weak' directives were added to the ISC source.

Category: Solaris-Specific

Action: added pragma weak directives

Files:

```
resolv/res_send.c
resolv/res_data.c
```

```
    resolv/res_debug.c
    resolv/res_comp.c
```

8. Add 'pragma weak' to retain __b64_ntop and __b64_pton functions.

Category: Solaris-Specific

Action: added pragma weak directives

Files:

```
    isc/base64.c
```

9. Some function prototypes in <resolv.h> use 'uint_t' (i.e. 'unsigned int'), while ISC's use 'u_long' (i.e. 'unsigned long').

Category: Solaris-Specific

Action: wrapped p_option and p_secstodate in sunw_wrappers.c so there are no changes in ISC code

Files:

```
    resolv/res_debug.c
```

10. Some function prototypes in <netdb.h> and <arpa/inet.h> use 'in_addr_t' while ISC's use 'u_long' or 'unsigned long.'

Category: Solaris-Specific

Action: renamed all functions to res_* in headers outside of ISC code.

Files:

```
    irs/getnetent.c
    inet/inet_lnaof.c
    inet/inet_makeaddr.c
    inet/inet_netof.c
    inet/inet_network.c
```

11. Some function prototypes in <arpa/inet.h> and <netdb.h> use 'socklen_t' while ISC's use 'size_t.'

Category: Solaris-Specific

Action: inet_ntop is not compiled. gethostent and getnameinfo renamed to res_*

Files:

```
    inet/inet_ntop.c
    irs/gethostent.c
    irs/getnameinfo.c
```

12. Get rid of "__func" #defines

Category: Solaris-Specific

Action: changes dropped. Updated private interface in mapfile-vers.

Files:

```
    isc/dst.h
    isc/tree.h
```

Structural Changes (and related implementation changes)

13. Add cylink, dnssafe, and irs subdirectories. Add EXPORT_SRC and CRYPT_SRC targets that delete all the cylink and dnssafe code.

Category: Solaris-Specific

Action: cylink and dnssafe removed. irs consolidated into libresolv2.

Files:

Makefile

14. Added copyright for the RSA Inc. files in the BIND distribution. Supercedes the copyright of the files in the dnssafe directory.

Category: Solaris-Specific

Action: dropped. dnssafe and cylink removed.

Files:

dnssafe/README.copyright

15. Build the /usr/lib/dns/cylink.so.1 and /usr/lib/dns/dnssafe.so.1 libraries. Add code to dynamically load the libraries (from /usr/lib/dns), and to execute the initialization functions for the respective encryption algorithms.

Category: Solaris-Specific

Action: dropped. dnssafe and cylink removed.

Files:

cylink/*

dnssafe/*

include/sunw_dst_init.h

sunw/sunw_dst_init.c

16. SC5.0 wants these constants to be 0x...U rather than 0x...L. (cylink is gone from ISC libbind)

Category: Solaris-Specific

Action: dropped. dnssafe and cylink removed.

Files:

cylink/sha.c

cylink/sha.h

17. Some additional code in dnssafe md5.h (dnssafe is gone from ISC libbind)

Category: Solaris-Specific

Action: dropped. dnssafe and cylink removed.

Files:

dnssafe/md5.h

18. Build the /usr/lib/dns/irs.so.1 library. Add code to dynamically load the /usr/lib/dns/irs.so.1 library, and to execute the initialization functions.

NOTE: The workaround in sunw_irs_init.c is to allow the functions irs_irp_acc and irs_nis_acc, which are not in libresolv.so.2 to be called from code within libresolv2.so.2. The __SUNW_IRS_INIT_NODEFINE macro tells the compiler not to perform this mapping when it is making irs.so.1.

Category: Solaris-Specific

Action: dropped. irs consolidated into libresolv2.

Files:

```
usr/src/lib/libresolv2/irs/*
usr/src/lib/libresolv2/include/sunw_irs_init.h
usr/src/lib/libresolv2/common/sunw/sunw_irs_init.c
```

Implementation Changes

19. Much of the private MT-safe resolver interface (see 1243174) used to live here. The BIND 8.2 resolver adds its own MT-safe interface; however, there are a few missing pieces to that MT-safety, so we need to retain parts of the thread-specific data (TSD) from the 1243174 fix, and to support Solaris binaries that used the old interface (DHCP dynamic DNS updates). Implement the 'set_no_hosts_fallback' and 'override_retry' fixes (see above as well as 4198434 and 4181371).

Category: ISC

Action: implemented functionality not yet implemented in ISC code

Files:

```
sunw/sunw_mtctxres.c
```

20. SUNW_INITCHKIF: Detect if there are no non-loopback interfaces so that we can fail quickly during boot when "dns" comes before "files" in nsswitch.conf. This fix was introduced a long time ago, and we carry on in the name of backward compatibility.

Category: ISC

Action: applied changes to resolv/res_init.c

Files:

```
resolv/res_init.c
```

21. SUNW_CONFCHECK and SUNW_AREWEINNAMED: Another hack to detect a broken configuration and avoid a timeout. SUNW_AREWEINNAMED is used to avoid triggering the SUNW_CONFCHECK code from inside the in.named.

Category: ISC

Action: Implemented but removed __areweinamed from common/mapfile/vers.

Files:

```
resolv/res_mkquery.c
```

22. SUNW_HOSTS_FALLBACK: When going through the NS switch, we must prevent libresolv from looking in /etc/hosts; 4198434.

Category: Solaris-Specific

Action: Applied changes

Files:

```
irs/gen.c
include/resolv_mt.h
sunw/sunw_mtctxres.c
```

23. SUNW_HNOK_UNDERSCORE: ISC's BIND disallows "_" in host names. While this is correct (they're illegal according to the relevant RFCs), Solaris has historically

allowed underscores in host names (and so did various old BIND versions), so we do too, in the name of consistency and backward compatibility.

Category: Solaris-Specific

Action: Applied changes

Files:

`resolv/res_comp.c`

24. SUNW_SETHERRNO: NS switch relies on `h_errno` always being set when there's an error. On Solaris, we must use the extended `ioctl`'s (`SIOCGLIFCONF` etc.) to get information for both IPv4 and IPv6 interfaces. Note that this code is used for the `libresolv` versions of `getipnodebyname/getipnodebyaddr`. While we don't expect customers to use them, they're useful during testing.

NOTE: although the functions where this fix is applied have been renamed because they overlap with `libnsl/libsocket` we'll still keep them for testing.

Category: ISC

Action: Applied changes

Files:

`irs/gethostent.c`

25. SUNW_OVERRIDE_RETRY: Surrounds code added to allow `nss_dns.so.1` to override `_res->retry` (the switch wants to set it to one and do its own NSL retries; see 4181371).

Category: ISC

Action: Applied changes

Files:

`irs/gethostent.c`
`include/resolv_mt.h`
`sunw/sunw_mtctxres.c`

26. SUNW_AVOIDSTDIO_FDLIMIT: Avoid 256 file descriptor limit in `stdio`

Category: ISC

Action: Dropped. This is not worth modifying the ISC code.

Files:

`isc/logging.c`
`irs/lcl_ho.c`
`resolv/res_init.c`
`resolv/res_query.c`

27. SUNW_LIBMD5: Use `md5(3EXT)` instead of internal implementation

NOTE: ISC implemented this as `HAVE_MD5`, but they no longer provide `dnssafe`, so that file would have to be changed.

Category: Solaris-specific (because `dnssafe` isn't part of ISC)

Action: Used new `HAVE_MD5` instead of `SUNW_LIBMD5`. Formerly `dst_hmac_md5_init()` was renamed to `dst_md5_hmac_init()`, compiled into `dnssafe.so.1` and loaded dynamically by `sunw_dst_hmac_md5_init()`. The wrapper, etc. was removed. Now we just compile it in and call it.

Files:

dst/hmac_link.c
dnssafe/md5.h

28. We need RFC 1535 disabled in order to retain implicit search paths (for backward compatibility, of course).

NOTE: no code changes were made to affect the behavior of “search” in /etc/resolv.conf

Category: Solaris-Specific

Action: No changes.

Files:

resolv/res_init.c

29. Undocumented changes: posix_getgrnam_r and posix_getpwnam_r have size_t instead of int for buflen:

NOTE: this looks like a cut and paste bug.

Category: N/A

Action: Dropped

Files:

irs/getgrent_r.c
irs/getpwent_r.c

30. Undocumented changes to irs_data.c. It appears that the diffs are because libresolv2 has an older version of this file.

Category: N/A

Action: Dropped

Files:

irs/irs_data.c

31. Add _IRS_DATA_H define so that it's safe to include this file multiple times.

Category: ISC

Action: Dropped. irs_data.h is private and not included by any other .h files

Files:

irs/irs_data.h

32. Silence compiler warning.

Category: N/A

Action: Dropped

Files:

isc/tree.c

33. Undocumented Change: add a function prototype.

Category: N/A

Action: Dropped. prototype is in isc/dst.h

Files:

nameser/ns_print.c

34. Undocumented change: functions `irs_destroy` and `net_data_minimize` removed from `irs.h`

Category: N/A

Action: Dropped. `net_data_minimize` was moved to `irs_data.h`

Files:

`include/irs.h`

35. Make sure `libresolv` itself doesn't refer to `__putlong` and `__putshort`. [4.4.1.1].

Category: Solaris-Specific

Action: Dropped. ISC code does not call `__putlong` or `__putshort`

Files:

`include/port_before.h`

`include/port_after.h`

36. `SUNW_OPTIONS`: Include the include file that contains all those other `SUNW_`-defines.

Category: Solaris-Specific

Action: Eliminated

```
SUNW_LIBNSL          /* conflicts for inet_addr, inet_ntoa */
SUNW_NSSEARCH        /* fix nslookup domain search */
SUNW_AVOIDOVERFLOW  /* Avoid buffer overflows */
SUNW_AVOIDOVERFLOW  /* Avoid 256 file descriptor limit in stdio */
SUNW_DOMAINFROMNIS  /* Default domain name from NIS/NIS+ */
USELOOPBACK          /* Resolver library defaults to 127.0.0.1 */
SUNW_AREWEINNAMED   /* Override _confcheck if proc is in.named */
SUNW_OPENFDOFFSET   /* Open non-stdio fd:s with offset */
SUNW_POLL            /* Use poll(2) instead of select(3) */
SUNW_LISTEN_BACKLOG /* Configurable listen(3N) backlog (named) */
SUNW_REJECT_BOGUS_H_LENGTH /* (libresolv) */
SUNW_QSTREAM_CLEANUP /* Avoid using free()d struct qstreams */
SUNW_CLOSEFROM      /* closefrom(3C) per PSARC 2000/193 */
```

Files:

`include/port_before.h`

Other changes:

37. removed useless `areweinnamed` function

38. updated public header files, removed `libbind` ones

39. added `port_*.h` files for public header workarounds

40. added workaround for `res_init_rnd`

41. removed `dnssafe/cylink` and updated `EXPORT_SRC` and `CRYPT_SRC` targets

42. removed obsolete `sunw_init` functions

43. fixed `sunw_dst_hmac_md5_init`

44. removed obsolete 'sunw' wrappers

45. removed/renamed overlapping irs functions, merged irs.so.1 into libresolv2
46. updated Makefile.com with dnssafe/cylink/irs changes
47. updated port_before.h, port_after.h, and sunoptions.h
48. extracted old libresolv2 changes into sunw_port_after.h
49. added May 2009 change to gethostent.c
50. added <sys/bitypes.h> to res_update.h
51. fixed AI_MASK bug
52. fixed sunw_uprec