

```

*****
7014 Sat Jan 5 00:24:56 2013
new/usr/src/cmd/sgs/libelf/common/getarsym.c
3451 archive libraries with no symbols shouldn't require a string table
*****
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) 1990, 2010, Oracle and/or its affiliates. All rights reserved.
24 */

26 /*      Copyright (c) 1988 AT&T */
27 /*      All Rights Reserved      */

29 #include <stdlib.h>
30 #include <errno.h>
31 #include <libelf.h>
32 #include "decl.h"
33 #include "msg.h"

36 /*
37  * Convert archive symbol table to memory format
38  *
39  * This takes a pointer to file's archive symbol table, alignment
40  * unconstrained. Returns null terminated vector of Elf_Arsym
41  * structures. Elf_Arsym uses size_t to represent offsets, which
42  * will be 32-bit in 32-bit versions, and 64-bits otherwise.
43  *
44  * There are two forms of archive symbol table, the original 32-bit
45  * form, and a 64-bit form originally found in IRIX64. The two formats
46  * differ only in the width of the integer word:
47  *
48  *      # offsets      4/8-byte word
49  *      offset[0...]  4/8-byte word each
50  *      strings        null-terminated, for offset[x]
51  *
52  * By default, the 64-bit form is only used when the archive exceeds
53  * the limits of 32-bits (4GB) in size. However, this is not required,
54  * and the ar -S option can be used to create a 64-bit symbol table in
55  * an archive that is under 4GB.
56  *
57  * Both 32 and 64-bit versions of libelf can read the 32-bit format
58  * without loss of information. Similarly, a 64-bit version of libelf
59  * will have no problem reading a 64-bit symbol table. This leaves the
60  * case where a 32-bit libelf reads a 64-bit symbol table, which requires
61  * some explanation. The offsets in a 64-bit symbol table will have zeros

```

```

62  * in the upper half of the words until the size of the archive exceeds 4GB.
63  * However, 32-bit libelf is unable to read any files larger than 2GB
64  * (see comments in update.c). As such, any archive that the 32-bit version
65  * of this code will encounter will be under 4GB in size. The upper 4
66  * bytes of each word will be zero, and can be safely ignored.
67  */

70 /*
71  * Offsets in archive headers are written in MSB (large endian) order
72  * on all platforms, regardless of native byte order. These macros read
73  * 4 and 8 byte values from unaligned memory.
74  *
75  * note:
76  * - The get8() macro for 32-bit code can ignore the first 4 bytes of
77  *   of the word, because they are known to be 0.
78  *
79  * - The inner most value in these macros is cast to an unsigned integer
80  *   of the final width in order to prevent the C compiler from doing
81  *   unwanted sign extension when the topmost bit of a byte is set.
82  */
83 #define get4(p) ((((((uint32_t)p[0]<<8)+p[1]<<8)+p[2]<<8)+p[3])
84
85 #ifdef _LP64
86 #define get8(p) (((((((((((uint64_t)p[0]<<8)+p[1]<<8)+p[2]<<8)+
87   p[3]<<8)+p[4]<<8)+p[5]<<8)+p[6]<<8)+p[7])
88 #else
89 #define get8(p) ((((((uint64_t)p[4]<<8)+p[5]<<8)+p[6]<<8)+p[7])
90 #endif

93 static Elf_Void *
94 arsym(Byte *off, size_t sz, size_t *e, int is64)
95 {
96     char          *endstr = (char *)off + sz;
97     register char *str;
98     Byte          *endoff;
99     Elf_Void      *oas;
100    size_t        eltsize = is64 ? 8 : 4;

102    {
103        register size_t n;

105        if (is64) {
106            if (sz < 8 || (sz - 8) / 8 < (n = get8(off))) {
107                _elf_seterr(EFMT_ARSYMSZ, 0);
108                return (NULL);
109                return (0);
110            } else {
111                if (sz < 4 || (sz - 4) / 4 < (n = get4(off))) {
112                    _elf_seterr(EFMT_ARSYMSZ, 0);
113                    return (NULL);
114                    return (0);
115                }
116                off += eltsize;
117                endoff = off + n * eltsize;

119                /*
120                 * If there are symbols in the symbol table, a
121                 * string table must be present and NULL terminated.
122                 *
123                 * The format dictates that the string table must always be
124                 * present, however in the case of an archive containing no
125                 * symbols GNU ar will not create one. We are permissive for

```

```

126      * the sake of compatibility.
127      * string table must be present, null terminated
128      */
129      if ((n > 0) && (((str = (char *)endoff) >= endstr) ||
130          (*(endstr - 1) != '\0'))) {
131
132          if (((str = (char *)endoff) >= endstr) ||
133              (*(endstr - 1) != '\0')) {
134              _elf_seterr(EFMT_ARSYM, 0);
135              return (NULL);
136          }
137          return (0);
138      }
139
140      /*
141      * There is always at least one entry returned if a symtab
142      * exists since the table's last entry is an artificial one
143      * with a NULL as_name, but is included in the count.
144      */
145      #endif /* ! codereview */
146      /*
147      * overflow can occur here, but not likely
148      */
149
150      *e = n + 1;
151      if ((oas = calloc(n + 1, sizeof (Elf_Arsym))) == NULL) {
152          n = sizeof (Elf_Arsym) * (n + 1);
153          if ((oas = malloc(n)) == 0) {
154              _elf_seterr(EMEM_ARSYM, errno);
155              return (NULL);
156          }
157          return (0);
158      }
159
160      register Elf_Arsym      *as = (Elf_Arsym *)oas;
161
162      while (off < endoff) {
163          if (str >= endstr) {
164              _elf_seterr(EFMT_ARSYMSTR, 0);
165              free(oas);
166              return (NULL);
167          }
168          if (is64)
169              as->as_off = get8(off);
170          else
171              as->as_off = get4(off);
172          as->as_name = str;
173          as->as_hash = elf_hash(str);
174          ++as;
175          off += eltsize;
176          while (*str++ != '\0')
177              /* LINTED */
178              ;
179      }
180      as->as_name = NULL;
181      as->as_name = 0;
182      as->as_off = 0;
183      as->as_hash = ~(unsigned long)0L;
184      return (oas);
185 }

```

unchanged_portion_omitted