

```

1  /* -- linux-c -*- ----- */
2  *
3  * Copyright (C) 1991, 1992 Linus Torvalds
4  * Copyright 2007 rPath, Inc. - All Rights Reserved
5  * Copyright 2009 Intel Corporation; author H. Peter Anvin
6  *
7  * This file is part of the Linux kernel, and is made available under
8  * the terms of the GNU General Public License version 2.
9  *
10 /*
11 */
12 /*
13 * Memory detection code
14 */
15
16 #include "boot.h"
17
18 #define SMAP 0x534d4150 /* ASCII "SMAP" */
19
20 static int detect_memory_e820(void)
21 {
22     int count = 0;
23     struct biosregs ireg, oreg;
24     struct e820entry *desc = boot_params.e820_map;
25     static struct e820entry buf; /* static so it is zeroed */
26
27     initregs(&ireg);
28     ireg.ax = 0xe820;
29     ireg.cx = sizeof buf;
30     ireg.edx = SMAP;
31     ireg.di = (size_t)&buf;
32
33     /*
34      * Note: at least one BIOS is known which assumes that the
35      * buffer pointed to by one e820 call is the same one as
36      * the previous call, and only changes modified fields. Therefore,
37      * we use a temporary buffer and copy the results entry by entry.
38      *
39      * This routine deliberately does not try to account for
40      * ACPI 3+ extended attributes. This is because there are
41      * BIOSes in the field which report zero for the valid bit for
42      * all ranges, and we don't currently make any use of the
43      * other attribute bits. Revisit this if we see the extended
44      * attribute bits deployed in a meaningful way in the future.
45     */
46
47     do {
48         intcall(0x15, &ireg, &oreg);
49         ireg.ebx = oreg.ebx; /* for next iteration... */
50
51         /* BIOSes which terminate the chain with CF = 1 as opposed
52          to %ebx = 0 don't always report the SMAP signature on
53          the final, failing, probe. */
54         if (oreg.eflags & X86_EFLAGS_CF)
55             break;
56
57         /* Some BIOSes stop returning SMAP in the middle of
58          the search loop. We don't know exactly how the BIOS
59          screwed up the map at that point, we might have a
60          partial map, the full map, or complete garbage, so
61          just return failure. */
62         if (oreg.eax != SMAP) {
63             count = 0;
64             break;
65         }
66
67         *desc++ = buf;
68         count++;
69     } while (ireg.ebx && count < ARRAY_SIZE(boot_params.e820_map));
70
71     return boot_params.e820_entries = count;
72 }
73

```

```

74 static int detect_memory_e801(void)
75 {
76     struct biosregs ireg, oreg;
77
78     initregs(&ireg);
79     ireg.ax = 0xe801;
80     intcall(0x15, &ireg, &oreg);
81
82     if (oreg.eflags & X86_EFLAGS_CF)
83         return -1;
84
85     /* Do we really need to do this? */
86     if (oreg.cx || oreg.dx) {
87         oreg.ax = oreg.cx;
88         oreg.bx = oreg.dx;
89     }
90
91     if (oreg.ax > 15*1024) {
92         return -1; /* Bogus! */
93     } else if (oreg.ax == 15*1024) {
94         boot_params.alt_mem_k = (oreg.dx << 6) + oreg.ax;
95     } else {
96         /*
97          * This ignores memory above 16MB if we have a memory
98          * hole there. If someone actually finds a machine
99          * with a memory hole at 16MB and no support for
100         * 0E820h they should probably generate a fake e820
101         * map.
102         */
103         boot_params.alt_mem_k = oreg.ax;
104     }
105
106     return 0;
107 }
108
109 static int detect_memory_88(void)
110 {
111     struct biosregs ireg, oreg;
112
113     initregs(&ireg);
114     ireg.ah = 0x88;
115     intcall(0x15, &ireg, &oreg);
116
117     boot_params.screen_info.ext_mem_k = oreg.ax;
118
119     return -(oreg.eflags & X86_EFLAGS_CF); /* 0 or -1 */
120 }
121
122 int detect_memory(void)
123 {
124     int err = -1;
125
126     if (detect_memory_e820() > 0)
127         err = 0;
128
129     if (!detect_memory_e801())
130         err = 0;
131
132     if (!detect_memory_88())
133         err = 0;
134
135     return err;
136 }
```