

new/usr/src/uts/i86pc/os/biosdisk.c

1

7961 Mon Mar 27 00:11:46 2017

new/usr/src/uts/i86pc/os/biosdisk.c

XXX nobios

```
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 */
21 /*
22 * Copyright 2008 Sun Microsystems, Inc. All rights reserved.
23 * Use is subject to license terms.
24 */
```

```
26 #include <sys/types.h>
27 #include <sys/param.h>
28 #include <sys/controlregs.h>
29 #include <sys/bootconf.h>
30 #include <sys/bootvfs.h>
31 #include <sys/bootregs.h>
32 #include <sys/bootconf.h>
33 #include <sys/conf.h>
34 #include <sys/promif.h>
35 #include <sys/ddi.h>
36 #include <sys/sunddi.h>
37 #include <sys/sunndi.h>
38 #include <sys/biosdisk.h>
39 #include <sys/psw.h>
40 #include <sys/machsystem.h>
41 #if defined(__xpv)
42 #include <sys/hypervisor.h>
43 #endif
44
45 extern int prom_debug;
46
47 /* hard code realmode memory address for now */
48 #define BIOS_RES_BUFFER_ADDR 0x7000
49
50 #define BIOSDEV_NUM 8
51 #define STARTING_DRVNUM 0x80
52 #define FP_OFF(fp) (((uintptr_t)(fp)) & 0xFFFF)
53 #define FP_SEG(fp) (((uintptr_t)(fp)) >> 16) & 0xFFFF
```

```
55 #ifdef DEBUG
56 int biosdebug = 0;
57 #define dprintf(fmt) \
58     if (biosdebug) \
59         prom_printf fmt
60 #else
61 #define dprintf(fmt)
```

new/usr/src/uts/i86pc/os/biosdisk.c

2

```
62 #endif
63
64 biosdev_data_t biosdev_info[BIOSDEV_NUM]; /* from 0x80 to 0x87 */
65 int dobiosdev = 1;
66
67
68 static int bios_check_extension_present(uchar_t);
69 static int get_dev_params(uchar_t);
70 static int read_firstblock(uchar_t drivenum);
71 static int drive_present(uchar_t drivenum);
72 static void reset_disk(uchar_t drivenum);
73 static int is_eltorito(uchar_t drivenum);
74
75 #if !defined(__xpv)
76 void
77 startup_bios_disk()
78 {
79     uchar_t drivenum;
80     int got_devparams = 0;
81     int got_first_block = 0;
82     uchar_t name[20];
83     dev_info_t *devi;
84     int extensions;
85
86     if (dobiosdev == 0 || !bios_calls_available) {
87         /*
88          * If BIOS calls have been disabled, or are not supported on
89          * this system, we cannot probe for the startup disk.
90          */
91         if (dobiosdev == 0)
92             return;
93
94         for (drivenum = 0x80; drivenum < (0x80 + BIOSDEV_NUM); drivenum++) {
95
96             if (!drive_present(drivenum))
97                 continue;
98
99             extensions = bios_check_extension_present(drivenum);
100
101             /*
102              * If we're booting from an Eltorito CD/DVD image, there's
103              * no need to get the device parameters or read the first block
104              * because we'll never install onto this device.
105              */
106             if (extensions && is_eltorito(drivenum))
107                 continue;
108
109             if (extensions && get_dev_params(drivenum))
110                 got_devparams = 1;
111             else
112                 got_devparams = 0;
113
114             if ((got_devparams == 1 || got_first_block == 0) &&
115                 /* retry */
116                 got_first_block = read_firstblock(drivenum);
117             }
118
119             if (got_devparams || got_first_block) {
120                 (void) sprintf((char *)name, "biosdev-0x%x", drivenum);
121                 devi = ddi_root_node();
122                 (void) e_ddi_prop_update_byte_array(DDI_DEV_T_NONE,
123                     devi, (char *)name,
124                     (uchar_t *)&biosdev_info[drivenum - 0x80],
125                     sizeof (biosdev_data_t));
126             }
127         }
128     }
129 }
```

new/usr/src/uts/i86pc/os/biosdisk.c

3

```
127     }  
128 }
```

_____unchanged_portion_omitted_

new/usr/src/uts/i86pc/os/fakebop.c

1

```
*****
63138 Mon Mar 27 00:11:47 2017
new/usr/src/uts/i86pc/os/fakebop.c
XXX nobios
*****
_____unchanged_portion_omitted_____

120 static bootprop_t *bprops = NULL;
121 static char *curr_page = NULL;          /* ptr to avail bprop memory */
122 static int curr_space = 0;             /* amount of memory at curr_page */

124 #ifdef __xpv
125 start_info_t *xen_info;
126 shared_info_t *HYPERVISOR_shared_info;
127 #endif

129 /*
130 * some allocator statistics
131 */
132 static ulong_t total_bop_alloc_scratch = 0;
133 static ulong_t total_bop_alloc_kernel = 0;

135 static void build_firmware_properties(void);

137 static int early_allocation = 1;

139 int force_fastreboot = 0;
140 volatile int fastreboot_onpanic = 0;
141 int post_fastreboot = 0;
142 #ifdef __xpv
143 volatile int fastreboot_capable = 0;
144 boolean_t bios_calls_available = B_FALSE;
145 #else
146 volatile int fastreboot_capable = 1;
147 boolean_t bios_calls_available = B_TRUE;
148 #endif

150 /*
151 * Information saved from current boot for fast reboot.
152 * If the information size exceeds what we have allocated, fast reboot
153 * will not be supported.
154 */
155 multiboot_info_t saved_mbi;
156 mb_memory_map_t saved_mmap[FASTBOOT_SAVED_MMMap_COUNT];
157 uint8_t saved_drives[FASTBOOT_SAVED_DRIVES_SIZE];
158 char saved_cmdline[FASTBOOT_SAVED_CMDLINE_LEN];
159 int saved_cmdline_len = 0;
160 size_t saved_file_size[FASTBOOT_MAX_FILES_MAP];

162 /*
163 * Turn off fastreboot_onpanic to avoid panic loop.
164 */
165 char fastreboot_onpanic_cmdline[FASTBOOT_SAVED_CMDLINE_LEN];
166 static const char fastreboot_onpanic_args[] = "-B fastreboot_onpanic=0";

168 /*
169 * Pointers to where System Resource Affinity Table (SRAT), System Locality
170 * Information Table (SLIT) and Maximum System Capability Table (MSCT)
171 * are mapped into virtual memory
172 */
173 ACPI_TABLE_SRAT *srat_ptr = NULL;
174 ACPI_TABLE_SLIT *slit_ptr = NULL;
175 ACPI_TABLE_MSCT *msct_ptr = NULL;

177 /*
178 * Arbitrary limit on number of localities we handle; if
```

new/usr/src/uts/i86pc/os/fakebop.c

2

```
179 * this limit is raised to more than UINT16_MAX, make sure
180 * process_slit() knows how to handle it.
181 */
182 #define SLIT_LOCALITIES_MAX (4096)

184 #define SLIT_NUM_PROPNAME "acpi-slit-localities"
185 #define SLIT_PROPNAME "acpi-slit"

187 /*
188 * Allocate aligned physical memory at boot time. This allocator allocates
189 * from the highest possible addresses. This avoids exhausting memory that
190 * would be useful for DMA buffers.
191 */
192 paddr_t
193 do_bop_phys_alloc(uint64_t size, uint64_t align)
194 {
195     paddr_t pa = 0;
196     paddr_t start;
197     paddr_t end;
198     struct memlist *ml = (struct memlist *)xbootp->bi_phys_install;

200     /*
201     * Be careful if high memory usage is limited in startup.c
202     * Since there are holes in the low part of the physical address
203     * space we can treat physmem as a pfn (not just a pgcnt) and
204     * get a conservative upper limit.
205     */
206     if (physmem != 0 && high_phys > pfn_to_pa(physmem))
207         high_phys = pfn_to_pa(physmem);

209     /*
210     * find the lowest or highest available memory in physinstalled
211     * On 32 bit avoid physmem above 4Gig if PAE isn't enabled
212     */
213 #if defined(__i386)
214     if (xbootp->bi_use_pae == 0 && high_phys > FOUR_GIG)
215         high_phys = FOUR_GIG;
216 #endif

218     /*
219     * find the highest available memory in physinstalled
220     */
221     size = P2ROUNDUP(size, align);
222     for (; ml; ml = ml->ml_next) {
223         start = P2ROUNDUP(ml->ml_address, align);
224         end = P2ALIGN(ml->ml_address + ml->ml_size, align);
225         if (start < next_phys)
226             start = P2ROUNDUP(next_phys, align);
227         if (end > high_phys)
228             end = P2ALIGN(high_phys, align);

230         if (end <= start)
231             continue;
232         if (end - start < size)
233             continue;

235     /*
236     * Early allocations need to use low memory, since
237     * physmem might be further limited by bootenv.rc
238     */
239     if (early_allocation) {
240         if (pa == 0 || start < pa)
241             pa = start;
242     } else {
243         if (end - size > pa)
244             pa = end - size;
```

new/usr/src/uts/i86pc/os/fakebop.c

```
245     }
246     }
247     if (pa != 0) {
248         if (early_allocation)
249             next_phys = pa + size;
250         else
251             high_phys = pa;
252         return (pa);
253     }
254     bop_panic("do_bop_phys_alloc(0x%" PRIx64 " , 0x%" PRIx64
255             ") Out of memory\n", size, align);
256     /*NOTREACHED*/
257 }
```

unchanged_portion_omitted

```
2514 /*
2515 * If this system has a PC-compatible BIOS, it will have handlers for
2516 * various well-known BIOS calls. These calls take the form of INT
2517 * instructions, revectoring to the nominated entry in the real mode
2518 * Interrupt Vector Table (IVT). If all of the commonly used entries (from
2519 * INT 10h up to INT 1Ah) are zero, we almost certainly don't want to make
2520 * use of BOP_DOINT() later.
2521 *
2522 * The IVT begins at linear address 0 on the 8086. Though later CPUs
2523 * allowed it to be moved, it seems that most BIOS implementations choose
2524 * not to do so for compatibility reasons. Our BIOS call trampoline (see
2525 * "idt_info" in "uts/i86pc/ml/bios_call_src.s") also assumes this address.
2526 */
2527 static int
2528 system_has_bios(void)
2529 {
2530     uint32_t all_ivts = 0;
2531
2532     DBG_MSG("\nBIOS IVT Entries:\n");
2533     for (uint32_t intnum = 0x10; intnum <= 0x1a; intnum++) {
2534         /*
2535          * The first software interrupt number (i.e. INT 0h) maps to
2536          * vector number 32 in the IVT. Each entry in the IVT is
2537          * four bytes, describing a 16 bit far call address.
2538          */
2539         uintptr_t slot = 4 * (32 + intnum);
2540         uint32_t ivte = *((uint32_t *)slot);
2541
2542         if (ivte != 0) {
2543             DBG(intnum);
2544             DBG(ivte);
2545         }
2546         all_ivts |= ivte;
2547     }
2548     if (all_ivts == 0) {
2549         DBG_MSG("System has no BIOS IVT entries\n");
2550     }
2551     DBG_MSG("\n");
2552
2553     return (all_ivts != 0);
2554 }
2555
2556 #else /* __xpv */
2557 static void
2558 enumerate_xen_cpus()
2559 {
2560     processorid_t id, max_id;
2561
2562     /*
2563      * User-set boot-ncpus overrides enumeration
2564      */

```

3

new/usr/src/uts/i86pc/os/fakebop.c

```
2565     if (do_bsys_getproplen(NULL, BOOT_NCPUS_NAME) >= 0)
2566         return;
2567
2568     /*
2569     * Probe every possible virtual CPU id and remember the
2570     * highest id present; the count of CPUs is one greater
2571     * than this. This tacitly assumes at least cpu 0 is present.
2572     */
2573     max_id = 0;
2574     for (id = 0; id < MAX_VIRT_CPUS; id++)
2575         if (HYPERVISOR_vcpu_op(VCPUOP_is_up, id, NULL) == 0)
2576             max_id = id;
2577
2578     bsetpropsi(BOOT_NCPUS_NAME, max_id+1);
2579
2580 }
2581 #endif /* __xpv */
2582
2583 static void
2584 build_firmware_properties(void)
2585 {
2586     ACPI_TABLE_HEADER *tp = NULL;
2587
2588     #ifndef __xpv
2589     if (do_bsys_getproplen(NULL, "no-bios") > 0 || !system_has_bios())
2590         bios_calls_available = B_FALSE;
2591
2592     if ((tp = find_fw_table(ACPI_SIG_MSCT)) != NULL)
2593         msct_ptr = process_msct((ACPI_TABLE_MSCT *)tp);
2594     else
2595         msct_ptr = NULL;
2596
2597     if ((tp = find_fw_table(ACPI_SIG_MADT)) != NULL)
2598         process_madt((ACPI_TABLE_MADT *)tp);
2599
2600     if ((srat_ptr = (ACPI_TABLE_SRAT *)
2601         find_fw_table(ACPI_SIG_SRAT)) != NULL)
2602         process_srat(srat_ptr);
2603
2604     if (slit_ptr = (ACPI_TABLE_SLIT *)find_fw_table(ACPI_SIG_SLIT))
2605         process_slit(slit_ptr);
2606
2607     tp = find_fw_table(ACPI_SIG_MCFG);
2608     #else /* __xpv */
2609     enumerate_xen_cpus();
2610     if (DOMAIN_IS_INITDOMAIN(xen_info))
2611         tp = find_fw_table(ACPI_SIG_MCFG);
2612     #endif /* __xpv */
2613     if (tp != NULL)
2614         process_mcfg((ACPI_TABLE_MCFG *)tp);
2615 }
2616
2617 unchanged_portion_omitted

```

4

new/usr/src/uts/i86pc/os/mlsetup.c

1

```
*****
14254 Mon Mar 27 00:11:48 2017
new/usr/src/uts/i86pc/os/mlsetup.c
XXX nobios
*****
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 */
21 /*
22  * Copyright (c) 2012 Gary Mills
23  *
24  * Copyright (c) 1993, 2010, Oracle and/or its affiliates. All rights reserved.
25  * Copyright (c) 2011 by Delphix. All rights reserved.
26  * Copyright 2016 Joyent, Inc.
27 */
28 /*
29  * Copyright (c) 2010, Intel Corporation.
30  * All rights reserved.
31 */

33 #include <sys/types.h>
34 #include <sys/sysmacros.h>
35 #include <sys/disp.h>
36 #include <sys/promif.h>
37 #include <sys/clock.h>
38 #include <sys/cpuvar.h>
39 #include <sys/stack.h>
40 #include <vm/as.h>
41 #include <vm/hat.h>
42 #include <sys/reboot.h>
43 #include <sys/avintr.h>
44 #include <sys/vtrace.h>
45 #include <sys/proc.h>
46 #include <sys/thread.h>
47 #include <sys/cpupart.h>
48 #include <sys/pset.h>
49 #include <sys/copyops.h>
50 #include <sys/pg.h>
51 #include <sys/disp.h>
52 #include <sys/debug.h>
53 #include <sys/sunddi.h>
54 #include <sys/x86_archext.h>
55 #include <sys/privregs.h>
56 #include <sys/machsystem.h>
57 #include <sys/onttrap.h>
58 #include <sys/bootconf.h>
59 #include <sys/boot_console.h>
60 #include <sys/kdi_machimpl.h>
61 #include <sys/archsystem.h>
```

new/usr/src/uts/i86pc/os/mlsetup.c

2

```
62 #include <sys/promif.h>
63 #include <sys/pci_cfgspace.h>
64 #include <sys/bootvfs.h>
65 #include <sys/tsc.h>
66 #ifdef __xpv
67 #include <sys/hypervisor.h>
68 #else
69 #include <sys/xpv_support.h>
70 #endif

72 /*
73  * some globals for patching the result of cpuid
74  * to solve problems w/ creative cpu vendors
75 */

77 extern uint32_t cpuid_feature_ecx_include;
78 extern uint32_t cpuid_feature_ecx_exclude;
79 extern uint32_t cpuid_feature_edx_include;
80 extern uint32_t cpuid_feature_edx_exclude;

82 /*
83  * Set console mode
84 */
85 static void
86 set_console_mode(uint8_t val)
87 {
88     struct bop_regs rp = {0};

90     if (!bios_calls_available)
91         return;

93     rp.eax.byte.ah = 0x0;
94     rp.eax.byte.al = val;
95     rp.ebx.word.bx = 0x0;

97     BOP_DOINT(bootops, 0x10, &rp);
98 }
_____unchanged_portion_omitted_____
```

```

*****
5862 Mon Mar 27 00:11:48 2017
new/usr/src/uts/i86pc/os/pci_bios.c
XXX nobios
*****
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 */
21 /*
22 * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All rights reserved.
23 */

25 #include <sys/types.h>
26 #include <sys/stat.h>
27 #include <sys/sunndi.h>
28 #include <sys/pci.h>
29 #include <sys/pci_impl.h>
30 #include <sys/pci_cfgspace.h>
31 #include <sys/pci_cfgspace_impl.h>
32 #include <sys/memlist.h>
33 #include <sys/bootconf.h>
34 #include <sys/psw.h>
35 #include <sys/machsystem.h>

37 /*
38  * pci irq routing information table
39  */
40 int          pci_irq_nroutes;
41 static pci_irq_route_t  *pci_irq_routes;

44 static int pci_bios_get_irq_routing(pci_irq_route_t *, int, int *);
45 static void pci_get_irq_routing_table(void);

48 /*
49  * Retrieve information from the bios needed for system
50  * configuration early during startup.
51  */
52 void
53 startup_pci_bios(void)
54 {
55     pci_get_irq_routing_table();
56 }

59 /*
60  * Issue the bios get irq routing information table interrupt
61  */

```

```

62  * Despite the name, the information in the table is only
63  * used to derive slot names for some named pci hot-plug slots.
64  *
65  * Returns the number of irq routing table entries returned
66  * by the bios, or 0 and optionally, the number of entries required.
67  */
68 static int
69 pci_bios_get_irq_routing(pci_irq_route_t *routes, int nroutes, int *nneededp)
70 {
71     struct bop_regs regs;
72     uchar_t      *hdrp;
73     uchar_t      *bufp;
74     int          i, n;
75     int          rval = 0;

77     if (nneededp)
78         *nneededp = 0;

80     /*
81      * If this system does not support BIOS calls, we can't use this
82      * mechanism.
83      */
84     if (!bios_calls_available)
85         return (0);

87     /*
88      * Set up irq routing header with the size and address
89      * of some useable low-memory data addresses.  Initialize
90      * data area to zero, avoiding memcpy/bzero.
91      */
92     hdrp = (uchar_t *)BIOS_IRQ_ROUTING_HDR;
93     bufp = (uchar_t *)BIOS_IRQ_ROUTING_DATA;

95     n = nroutes * sizeof (pci_irq_route_t);
96     for (i = 0; i < n; i++)
97         bufp[i] = 0;
98     ((pci_irq_route_hdr_t *)hdrp)->pir_size = n;
99     ((pci_irq_route_hdr_t *)hdrp)->pir_addr = (uint32_t)(uintptr_t)bufp;

101     bzero(&regs, sizeof (regs));
102     regs.eax.word.ax = (PCI_FUNCTION_ID << 8) | PCI_GET_IRQ_ROUTING;

104     regs.ds = 0xf000;
105     regs.es = FP_SEG((uint_t)(uintptr_t)hdrp);
106     regs.edi.word.di = FP_OFF((uint_t)(uintptr_t)hdrp);

108     BOP_DOINT(bootops, 0x1a, &regs);

110     n = (int)(((pci_irq_route_hdr_t *)hdrp)->pir_size /
111             sizeof (pci_irq_route_t));

113     if ((regs.eflags & PS_C) != 0) {
114         if (nneededp)
115             *nneededp = n;
116     } else {
117         /*
118          * Copy resulting irq routing data from low memory up to
119          * the kernel address space, avoiding memcpy as usual.
120          */
121         if (n <= nroutes) {
122             for (i = 0; i < n * sizeof (pci_irq_route_t); i++)
123                 ((uchar_t *)routes)[i] = bufp[i];
124             rval = n;
125         }
126     }
127     return (rval);

```

new/usr/src/uts/i86pc/os/pci_bios.c

3

128 }

unchanged_portion_omitted

new/usr/src/uts/i86pc/os/pci_cfgspace.c

1

```
*****
8846 Mon Mar 27 00:11:49 2017
new/usr/src/uts/i86pc/os/pci_cfgspace.c
XXX nobios
*****
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) 2005, 2010, Oracle and/or its affiliates. All rights reserved.
24 */

26 /*
27  * PCI configuration space access routines
28 */

30 #include <sys/system.h>
31 #include <sys/psw.h>
32 #include <sys/bootconf.h>
33 #include <sys/reboot.h>
34 #include <sys/pci_impl.h>
35 #include <sys/pci_cfgspace.h>
36 #include <sys/pci_cfgspace_impl.h>
37 #include <sys/pci_cfgacc.h>
38 #include <sys/machsystem.h>
39 #if defined(__xpv)
40 #include <sys/hypervisor.h>
41 #endif

42 #if defined(__xpv)
43 int pci_max_nbus = 0xFE;
44 #endif
44 int pci_bios_cfg_type = PCI_MECHANISM_UNKNOWN;
45 int pci_bios_maxbus;
46 int pci_bios_mech;
47 int pci_bios_vers;

49 /*
50  * These two variables can be used to force a configuration mechanism or
51  * to force which function is used to probe for the presence of the PCI bus.
52  */
53 int PCI_CFG_TYPE = 0;
54 int PCI_PROBE_TYPE = 0;

56 /*
57  * No valid mcfg_mem_base by default, and accessing pci config space
58  * in mem-mapped way is disabled.
59  */
```

new/usr/src/uts/i86pc/os/pci_cfgspace.c

2

```
60 uint64_t mcfg_mem_base = 0;
61 uint8_t mcfg_bus_start = 0;
62 uint8_t mcfg_bus_end = 0xFF;

64 /*
65  * Maximum offset in config space when not using MMIO
66 */
67 uint_t pci_iocfg_max_offset = 0xFF;

69 /*
70  * These function pointers lead to the actual implementation routines
71  * for configuration space access. Normally they lead to either the
72  * pci_mech1_* or pci_mech2_* routines, but they can also lead to
73  * routines that work around chipset bugs.
74  * These functions are accessing pci config space via I/O way.
75  * Pci_cfgacc_get/put functions should be used as more common interfaces,
76  * which also provide accessing pci config space via mem-mapped way.
77 */
78 uint8_t (*pci_getb_func)(int bus, int dev, int func, int reg);
79 uint16_t (*pci_getw_func)(int bus, int dev, int func, int reg);
80 uint32_t (*pci_getl_func)(int bus, int dev, int func, int reg);
81 void (*pci_putb_func)(int bus, int dev, int func, int reg, uint8_t val);
82 void (*pci_putw_func)(int bus, int dev, int func, int reg, uint16_t val);
83 void (*pci_putl_func)(int bus, int dev, int func, int reg, uint32_t val);

85 extern void (*pci_cfgacc_acc_p)(pci_cfgacc_req_t *req);

87 /*
88  * Internal routines
89 */
90 static int pci_check(void);

92 #if !defined(__xpv)
93 static int pci_check_bios(void);
94 static int pci_get_cfg_type(void);
95 #endif

97 /* for legacy io-based config space access */
98 kmutex_t pcicfg_mutex;

100 /* for mmio-based config space access */
101 kmutex_t pcicfg_mmio_mutex;

103 /* ..except Orion and Neptune, which have to have their own */
104 kmutex_t pcicfg_chipset_mutex;

106 void
107 pci_cfgspace_init(void)
108 {
109     mutex_init(&pcicfg_mutex, NULL, MUTEX_SPIN,
110         (ddi_iblock_cookie_t)ipltospl(15));
111     mutex_init(&pcicfg_mmio_mutex, NULL, MUTEX_SPIN,
112         (ddi_iblock_cookie_t)ipltospl(DISPLAY_LEVEL));
113     mutex_init(&pcicfg_chipset_mutex, NULL, MUTEX_SPIN,
114         (ddi_iblock_cookie_t)ipltospl(15));
115     if (!pci_check()) {
116         mutex_destroy(&pcicfg_mutex);
117         mutex_destroy(&pcicfg_mmio_mutex);
118         mutex_destroy(&pcicfg_chipset_mutex);
119     }
120 }

    unchanged_portion_omitted

246 #if !defined(__xpv)

248 static int
```



```
249 pci_check_bios(void)
250 {
251     struct bop_regs regs;
252     uint32_t    carryflag;
253     uint16_t    ax, dx;
254
255     if (!bios_calls_available) {
256         /*
257          * If this system does not support BIOS calls, we must fall
258          * back to default values and a search of all of the possible
259          * PCI buses.
260          */
261         pci_bios_mech = 1;
262         pci_bios_vers = 0;
263         pci_bios_maxbus = pci_max_nbus;
264         return (PCI_MECHANISM_1);
265     }
266
267     bzero(&regs, sizeof (regs));
268     regs.eax.word.ax = (PCI_FUNCTION_ID << 8) | PCI_BIOS_PRESENT;
269
270     BOP_DOINT(bootops, 0x1a, &regs);
271     carryflag = regs.eflags & PS_C;
272     ax = regs.eax.word.ax;
273     dx = regs.edx.word.dx;
274
275     /* the carry flag must not be set */
276     if (carryflag != 0)
277         return (PCI_MECHANISM_NONE);
278
279     if (dx != ('P' | 'C' << 8))
280         return (PCI_MECHANISM_NONE);
281
282     /* ah (the high byte of ax) must be zero */
283     if ((ax & 0xff00) != 0)
284         return (PCI_MECHANISM_NONE);
285
286     pci_bios_mech = (ax & 0x3);
287     pci_bios_vers = regs.ebx.word.bx;
288     pci_bios_maxbus = (regs.ecx.word.cx & 0xff);
289
290     switch (pci_bios_mech) {
291     default: /* ??? */
292     case 0: /* supports neither? */
293         return (PCI_MECHANISM_NONE);
294
295     case 1:
296     case 3: /* supports both */
297         return (PCI_MECHANISM_1);
298
299     case 2:
300         return (PCI_MECHANISM_2);
301     }
302 }
```

unchanged portion omitted

new/usr/src/uts/i86pc/sys/machsystem.h

1

```
*****
6588 Mon Mar 27 00:11:49 2017
new/usr/src/uts/i86pc/sys/machsystem.h
XXX nobios
*****
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) 1993, 2010, Oracle and/or its affiliates. All rights reserved.
24 */
25 /*
26  * Copyright (c) 2010, Intel Corporation.
27  * All rights reserved.
28 */

30 #ifndef _SYS_MACHSYSTEM_H
31 #define _SYS_MACHSYSTEM_H

33 /*
34  * Numerous platform-dependent interfaces that don't seem to belong
35  * in any other header file.
36  *
37  * This file should not be included by code that purports to be
38  * platform-independent.
39 */

41 #include <sys/machparam.h>
42 #include <sys/varargs.h>
43 #include <sys/thread.h>
44 #include <sys/cpuvar.h>
45 #include <sys/privregs.h>
46 #include <sys/system.h>
47 #include <sys/traptrace.h>
48 #include <vm/page.h>

50 #ifdef __cplusplus
51 extern "C" {
52 #endif

54 #ifdef _KERNEL

56 typedef enum mach_cpu_add_arg_type {
57     MACH_CPU_ARG_LOCAL_APIC,
58     MACH_CPU_ARG_LOCAL_X2APIC,
59 } mach_cpu_add_arg_type_t;
    unchanged_portion_omitted_

```

new/usr/src/uts/i86pc/sys/machsystem.h

2

```
224 /* Maximum physical page number (PFN) for memory DR operations. */
225 extern uint64_t plat_dr_physmax;

227 #ifdef __xpv
228 #include <sys/xen_mmu.h>
229 extern page_t *page_get_high_mfn(mfn_t);
230 #endif

232 extern hrtime_t tsc_gethrtime_tick_delta(void);

234 extern boolean_t bios_calls_available;

236 #endif /* _KERNEL */

238 #ifdef __cplusplus
239 }
    unchanged_portion_omitted_

```