

new/usr/src/uts/i86pc/apix/Makefile

1

```
*****
2037 Thu Jan 17 14:22:52 2019
new/usr/src/uts/i86pc/apix/Makefile
10246 fix apic_allocate_irq() indentation
*****
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 # uts/i86pc/apix/Makefile
23 #
24 # Copyright (c) 2010, Oracle and/or its affiliates. All rights reserved.
25 # Copyright (c) 2019, Joyent, Inc.
25 # Copyright (c) 2018, Joyent, Inc.
26 #
27 # This makefile drives the production of the apix "mach"
28 # kernel module.
29 #
30 # apix implementation architecture dependent
31 #
32 #
33 #
34 # Path to the base of the uts directory tree (usually /usr/src/uts).
35 #
36 UTSBASE = ../../

38 #
39 # Define the module and object file sets.
40 #
41 MODULE = apix
42 OBJECTS = ${APIX_OBJS:%=${OBJS_DIR}/%}
43 LINTS = ${APIX_OBJS:%.o=${LINTS_DIR}/%.ln}
44 ROOTMODULE = ${ROOT_PSM_MACH_DIR}/${MODULE}

46 #
47 # Include common rules.
48 #
49 include ${UTSBASE}/i86pc/Makefile.i86pc

51 #
52 # Define targets
53 #
54 ALL_TARGET = ${BINARY}
55 LINT_TARGET = ${MODULE}.lint
56 INSTALL_TARGET = ${BINARY} ${ROOTMODULE}

58 DEBUG_FLGS =
59 $(NOT_RELEASE_BUILD)DEBUG_DEFS += ${DEBUG_FLGS}
```

new/usr/src/uts/i86pc/apix/Makefile

2

```
61 #
62 # Depends on ACPI CA interpreter
63 #
64 LDFLAGS += -dy -N misc/acpica

66 # needs work, real bug in apic_allocate_irq()
67 ${OBJS_DIR}/mp_platform_common.o := SMOFF += indenting
68 # needs work
69 ${OBJS_DIR}/psm_common.o := SMOFF += deref_check

69 #
70 # Default build targets.
71 #
72 .KEEP_STATE:

74 def: $(DEF_DEPS)

76 all: $(ALL_DEPS)

78 clean: $(CLEAN_DEPS)

80 clobber: $(CLOBBER_DEPS)

82 lint: $(LINT_DEPS)

84 modlintlib: $(MODLINTLIB_DEPS)

86 clean.lint: $(CLEAN_LINT_DEPS)

88 install: $(INSTALL_DEPS)

90 #
91 # Include common targets.
92 #
93 include ${UTSBASE}/i86pc/Makefile.targ
```

```

*****
68319 Thu Jan 17 14:22:52 2019
new/usr/src/uts/i86pc/io/mp_platform_common.c
10246 fix apic_allocate_irq() indentation
*****
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) 2007, 2010, Oracle and/or its affiliates. All rights reserved.
23 * Copyright 2016 Nexenta Systems, Inc.
24 * Copyright (c) 2017 by Delphix. All rights reserved.
25 * Copyright (c) 2019, Joyent, Inc.
26 * Copyright (c) 2018, Joyent, Inc.
27 */
28 * Copyright (c) 2010, Intel Corporation.
29 * All rights reserved.
30 */

32 /*
33 * PSMI 1.1 extensions are supported only in 2.6 and later versions.
34 * PSMI 1.2 extensions are supported only in 2.7 and later versions.
35 * PSMI 1.3 and 1.4 extensions are supported in Solaris 10.
36 * PSMI 1.5 extensions are supported in Solaris Nevada.
37 * PSMI 1.6 extensions are supported in Solaris Nevada.
38 * PSMI 1.7 extensions are supported in Solaris Nevada.
39 */
40 #define PSMI_1_7

42 #include <sys/processor.h>
43 #include <sys/time.h>
44 #include <sys/psm.h>
45 #include <sys/smp_impldefs.h>
46 #include <sys/cram.h>
47 #include <sys/acpi/acpi.h>
48 #include <sys/acpica.h>
49 #include <sys/psm_common.h>
50 #include <sys/apic.h>
51 #include <sys/apic_timer.h>
52 #include <sys/pit.h>
53 #include <sys/ddi.h>
54 #include <sys/sunndi.h>
55 #include <sys/ddi_impldefs.h>
56 #include <sys/pci.h>
57 #include <sys/promif.h>
58 #include <sys/x86_archext.h>
59 #include <sys/cpc_impl.h>
60 #include <sys/uadmin.h>

```

```

61 #include <sys/panic.h>
62 #include <sys/debug.h>
63 #include <sys/archsystem.h>
64 #include <sys/trap.h>
65 #include <sys/machsystem.h>
66 #include <sys/cpuvar.h>
67 #include <sys/rm_platter.h>
68 #include <sys/privregs.h>
69 #include <sys/cyclic.h>
70 #include <sys/note.h>
71 #include <sys/pci_intr_lib.h>
72 #include <sys/sunndi.h>
73 #if !defined(__xpv)
74 #include <sys/hpet.h>
75 #include <sys/clock.h>
76 #endif

78 /*
79  * Local Function Prototypes
80 */
81 static int apic_handle_defconf();
82 static int apic_parse_mpct(caddr_t mpct, int bypass);
83 static struct apic_mpfps_hdr *apic_find_fps_sig(caddr_t fptr, int size);
84 static int apic_checksum(caddr_t bptr, int len);
85 static int apic_find_bus_type(char *bus);
86 static int apic_find_bus(int busid);
87 static struct apic_io_intr *apic_find_io_intr(int irqno);
88 static int apic_find_free_irq(int start, int end);
89 struct apic_io_intr *apic_find_io_intr_w_busid(int irqno, int busid);
90 static void apic_set_pwroff_method_from_mpcnfhdr(struct apic_mp_cnf_hdr *hdrp);
91 static void apic_free_apic_cpus(void);
92 static boolean_t apic_is_ioapic_AMD_813x(uint32_t physaddr);
93 static int apic_acpi_enter_apicmode(void);

95 int apic_handle_pci_pci_bridge(dev_info_t *dip, int child_devno,
96 int child_ipin, struct apic_io_intr **intrp);
97 int apic_find_bus_id(int bustype);
98 int apic_find_intin(uchar_t ioapic, uchar_t intin);
99 void apic_record_rdt_entry(apic_irq_t *irqptr, int irq);

101 int apic_debug_mps_id = 0; /* 1 - print MPS ID strings */

103 /* ACPI SCI interrupt configuration; -1 if SCI not used */
104 int apic_sci_vect = -1;
105 iflag_t apic_sci_flags;

107 #if !defined(__xpv)
108 /* ACPI HPET interrupt configuration; -1 if HPET not used */
109 int apic_hpet_vect = -1;
110 iflag_t apic_hpet_flags;
111 #endif

113 /*
114  * psm name pointer
115 */
116 char *psm_name;

118 /* ACPI support routines */
119 static int acpi_probe(char *);
120 static int apic_acpi_irq_configure(acpi_psm_lnk_t *acpipsm_lnkp, dev_info_t *dip,
121 int *pci_irqp, iflag_t *intr_flagp);

123 int apic_acpi_translate_pci_irq(dev_info_t *dip, int busid, int devid,
124 int ipin, int *pci_irqp, iflag_t *intr_flagp);
125 uchar_t apic_find_ioapic(int irq);
126 static int acpi_intr_compatible(iflag_t iflag1, iflag_t iflag2);

```

```

128 /* Max wait time (in repetitions) for flags to clear in an RDT entry. */
129 int apic_max_reps_clear_pending = 1000;

131 int      apic_intr_policy = INTR_ROUND_ROBIN;

133 int      apic_next_bind_cpu = 1; /* For round robin assignment */
134          /* start with cpu 1 */

136 /*
137 * If enabled, the distribution works as follows:
138 * On every interrupt entry, the current ipl for the CPU is set in cpu_info
139 * and the irq corresponding to the ipl is also set in the aci_current array.
140 * interrupt exit and setspl (due to soft interrupts) will cause the current
141 * ipl to be changed. This is cache friendly as these frequently used
142 * paths write into a per cpu structure.
143 *
144 * Sampling is done by checking the structures for all CPUs and incrementing
145 * the busy field of the irq (if any) executing on each CPU and the busy field
146 * of the corresponding CPU.
147 * In periodic mode this is done on every clock interrupt.
148 * In one-shot mode, this is done thru a cyclic with an interval of
149 * apic_redistribute_sample_interval (default 10 milli sec).
150 *
151 * Every apic_sample_factor_redistribution times we sample, we do computations
152 * to decide which interrupt needs to be migrated (see comments
153 * before apic_intr_redistribute()).
154 */

156 /*
157 * Following 3 variables start as % and can be patched or set using an
158 * API to be defined in future. They will be scaled to
159 * sample_factor_redistribution which is in turn set to hertz+1 (in periodic
160 * mode), or 101 in one-shot mode to stagger it away from one sec processing
161 */

163 int      apic_int_busy_mark = 60;
164 int      apic_int_free_mark = 20;
165 int      apic_diff_for_redistribution = 10;

167 /* sampling interval for interrupt redistribution for dynamic migration */
168 int      apic_redistribute_sample_interval = NANOSEC / 100; /* 10 millisecc */

170 /*
171 * number of times we sample before deciding to redistribute interrupts
172 * for dynamic migration
173 */
174 int      apic_sample_factor_redistribution = 101;

176 int      apic_redist_cpu_skip = 0;
177 int      apic_num_imbalance = 0;
178 int      apic_num_rebind = 0;

180 /*
181 * Maximum number of APIC CPUs in the system, -1 indicates that dynamic
182 * allocation of CPU ids is disabled.
183 */
184 int      apic_max_nproc = -1;
185 int      apic_nproc = 0;
186 size_t   apic_cpus_size = 0;
187 int      apic_defconf = 0;
188 int      apic_irq_translate = 0;
189 int      apic_spec_rev = 0;
190 int      apic_imcrp = 0;

192 int      apic_use_acpi = 1;      /* 1 = use ACPI, 0 = don't use ACPI */

```

```

193 int      apic_use_acpi_madt_only = 0;      /* 1=ONLY use MADT from ACPI */

195 /*
196 * For interrupt link devices, if apic_unconditional_srs is set, an irq resource
197 * will be assigned (via _SRS). If it is not set, use the current
198 * irq setting (via _CRS), but only if that irq is in the set of possible
199 * irqs (returned by _PRS) for the device.
200 */
201 int      apic_unconditional_srs = 1;

203 /*
204 * For interrupt link devices, if apic_prefer_crs is set when we are
205 * assigning an IRQ resource to a device, prefer the current IRQ setting
206 * over other possible irq settings under same conditions.
207 */

209 int      apic_prefer_crs = 1;

211 uchar_t  apic_io_id[MAX_IO_APIC];
212 volatile uint32_t *apicioadr[MAX_IO_APIC];
213 uchar_t  apic_io_ver[MAX_IO_APIC];
214 uchar_t  apic_io_vectbase[MAX_IO_APIC];
215 uchar_t  apic_io_vectend[MAX_IO_APIC];
216 uchar_t  apic_reserved_irqlist[MAX_ISA_IRQ + 1];
217 uint32_t apic_physaddr[MAX_IO_APIC];

219 boolean_t ioapic_mask_workaround[MAX_IO_APIC];

221 /*
222 * First available slot to be used as IRQ index into the apic_irq_table
223 * for those interrupts (like MSI/X) that don't have a physical IRQ.
224 */
225 int      apic_first_avail_irq = APIC_FIRST_FREE_IRQ;

227 /*
228 * apic_ioapic_lock protects the ioapics (reg select), the status, temp_bound
229 * and bound elements of cpus_info and the temp_cpu element of irq_struct
230 */
231 lock_t   apic_ioapic_lock;

233 int      apic_io_max = 0;      /* no. of i/o apics enabled */

235 struct  apic_io_intr *apic_io_intrp = NULL;
236 static struct  apic_bus *apic_busp;

238 uchar_t  apic_resv_vector[MAXIPL+1];

240 char     apic_level_intr[APIC_MAX_VECTOR+1];

242 uint32_t      eisa_level_intr_mask = 0;
243          /* At least MSB will be set if EISA bus */

245 int      apic_pci_bus_total = 0;
246 uchar_t  apic_single_pci_busid = 0;

248 /*
249 * airq_mutex protects additions to the apic_irq_table - the first
250 * pointer and any airq_nexts off of that one. It also protects
251 * apic_max_device_irq & apic_min_device_irq. It also guarantees
252 * that share_id is unique as new ids are generated only when new
253 * irq_t structs are linked in. Once linked in the structs are never
254 * deleted. temp_cpu & mps_intr_index field indicate if it is programmed
255 * or allocated. Note that there is a slight gap between allocating in
256 * apic_intrap_xlate and programming in addspl.
257 */
258 kmutex_t      airq_mutex;

```

```

259 apic_irq_t      *apic_irq_table[APIC_MAX_VECTOR+1];
260 int             apic_max_device_irq = 0;
261 int             apic_min_device_irq = APIC_MAX_VECTOR;

263 typedef struct prs_irq_list_ent {
264     int          list_prio;
265     int32_t      irq;
266     iflag_t      intrflags;
267     acpi_prs_private_t prsprv;
268     struct prs_irq_list_ent *next;
269 } prs_irq_list_t;
    unchanged_portion_omitted

1602 int
1603 apic_allocate_irq(int irq)
1604 {
1605     int     freeirq, i;

1607     if ((freeirq = apic_find_free_irq(irq, (APIC_RESV_IRQ - 1))) == -1) {
1607         if ((freeirq = apic_find_free_irq(irq, (APIC_RESV_IRQ - 1))) == -1)
1608             if ((freeirq = apic_find_free_irq(APIC_FIRST_FREE_IRQ,
1609                 (irq - 1))) == -1) {
1610                 /*
1611                  * if BIOS really defines every single irq in the mps
1612                  * table, then don't worry about conflicting with
1613                  * them, just use any free slot in apic_irq_table
1614                  */
1615                 for (i = APIC_FIRST_FREE_IRQ; i < APIC_RESV_IRQ; i++) {
1616                     if ((apic_irq_table[i] == NULL) ||
1617                         apic_irq_table[i]->airq_mps_intr_index ==
1618                             FREE_INDEX) {
1619                         freeirq = i;
1620                         break;
1621                     }
1622                 }

1624                 if (freeirq == -1) {
1625                     /* This shouldn't happen, but just in case */
1626                     cmn_err(CE_WARN, "%s: NO available IRQ", psm_name);
1627                     return (-1);
1628                 }
1629             }
1630     }

1632     if (apic_irq_table[freeirq] == NULL) {
1633         apic_irq_table[freeirq] =
1634             kmem_zalloc(sizeof (apic_irq_t), KM_NOSLEEP);
1635         if (apic_irq_table[freeirq] == NULL) {
1636             cmn_err(CE_WARN, "%s: NO memory to allocate IRQ",
1637                 psm_name);
1638             return (-1);
1639         }
1640         apic_irq_table[freeirq]->airq_temp_cpu = IRQ_UNINIT;
1641         apic_irq_table[freeirq]->airq_mps_intr_index = FREE_INDEX;
1642     }
1643     return (freeirq);
1644 }
    unchanged_portion_omitted

```

new/usr/src/uts/i86pc/pcplusmp/Makefile

1

```
*****
2090 Thu Jan 17 14:22:53 2019
new/usr/src/uts/i86pc/pcplusmp/Makefile
10246 fix apic_allocate_irq() indentation
*****
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 # uts/i86pc/pcplusmp/Makefile
23 #
24 # Copyright 2007 Sun Microsystems, Inc. All rights reserved.
25 # Use is subject to license terms.
26 #
27 # Copyright (c) 2019, Joyent, Inc.
27 # Copyright (c) 2018, Joyent, Inc.
28 #

30 #
31 # This makefile drives the production of the pcplusmp "mach"
32 # kernel module.
33 #
34 # pcplusmp implementation architecture dependent
35 #

37 #
38 # Path to the base of the uts directory tree (usually /usr/src/uts).
39 #
40 UTSBASE = ../../

42 #
43 # Define the module and object file sets.
44 #
45 MODULE = pcplusmp
46 OBJECTS = $(PCPLUSMP_OBJS:%=$(OBJS_DIR)/%)
47 LINTS = $(PCPLUSMP_OBJS:%.o=$(LINTS_DIR)/%.ln)
48 ROOTMODULE = $(ROOT_PSM_MACH_DIR)/$(MODULE)

50 #
51 # Include common rules.
52 #
53 include $(UTSBASE)/i86pc/Makefile.i86pc

55 #
56 # Define targets
57 #
58 ALL_TARGET = $(BINARY)
59 LINT_TARGET = $(MODULE).lint
60 INSTALL_TARGET = $(BINARY) $(ROOTMODULE)
```

new/usr/src/uts/i86pc/pcplusmp/Makefile

2

```
62 DEBUG_FLGS =
63 $(NOT_RELEASE_BUILD)DEBUG_DEFS += $(DEBUG_FLGS)

65 #
66 # Depends on ACPI CA interpreter
67 #
68 LDFLAGS += -dy -N misc/acpica

70 # needs work
71 $(OBJS_DIR)/mp_platform_common.o := SMOFF += indenting
72 # needs work
71 $(OBJS_DIR)/psm_common.o := SMOFF += deref_check

73 #
74 # Default build targets.
75 #
76 .KEEP_STATE:

78 def: $(DEF_DEPS)

80 all: $(ALL_DEPS)

82 clean: $(CLEAN_DEPS)

84 clobber: $(CLOBBER_DEPS)

86 lint: $(LINT_DEPS)

88 modlintlib: $(MODLINTLIB_DEPS)

90 clean.lint: $(CLEAN_LINT_DEPS)

92 install: $(INSTALL_DEPS)

94 #
95 # Include common targets.
96 #
97 include $(UTSBASE)/i86pc/Makefile.targ
```