

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

```
*****
1978 Thu Sep 7 15:25:31 2017
new/usr/src/uts/i86pc/apix/Makefile
8626 make pcplusmp and apix warning-free
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
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 2016, Joyent, Inc.
26 #
27 # This makefile drives the production of the pcplusmp "mach"
28 # kernel module.
29 #
30 # pcplusmp implementation architecture dependent
31 #

32 #
33 # Path to the base of the uts directory tree (usually /usr/src/uts).
34 #
35 #
36 UTSBASE = ../../

37 #
38 # Define the module and object file sets.
39 #
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)

45 #
46 # Include common rules.
47 #
48 #
49 include $(UTSBASE)/i86pc/Makefile.i86pc

50 #
51 # Define targets
52 #
53 #
54 ALL_TARGET   = $(BINARIES)
55 LINT_TARGET  = $(MODULE).lint
56 INSTALL_TARGET = $(BINARIES) $(ROOTMODULE)

57 DEBUG_FLGS   =
58 $(NOT_RELEASE_BUILD)DEBUG_DEFS += $(DEBUG_FLGS)
```

1

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

```
61 #
62 # Depends on ACPI CA interpreter
63 #
64 LDFLAGS      += -dy -N misc/acpica
65 CERRWARN    += -_gcc=-Wno-uninitialized
66 CERRWARN    += -_gcc=-Wno-unused-function
67 #
68 # Default build targets.
69 .KEEP_STATE:

70 def:          $(DEF_DEPS)
71 all:          $(ALL_DEPS)
72 clean:        $(CLEAN_DEPS)
73 clobber:     $(CLOBBER_DEPS)
74 lint:         $(LINT_DEPS)
75 modlintlib:  $(MODLINTLIB_DEPS)
76 clean.lint:   $(CLEAN_LINT_DEPS)
77 install:     $(INSTALL_DEPS)

78 #
79 # Include common targets.
80 #
81 include $(UTSBASE)/i86pc/Makefile.targ
```

2


```
1719             (void *)APIX_GET_DIP(vp)));
1720     }
1721 #endif /* DEBUG */
1722 }
1723
1725     cap_ptr = i_ddi_get_msi_msix_cap_ptr(dip);
1726     handle = i_ddi_get_pci_config_handle(dip);
1727     msi_ctrl = pci_config_get16(handle, cap_ptr + PCI_MSI_CTRL);
1728
1729     /* MSI Per vector masking is supported. */
1730     if (msi_ctrl & PCI_MSI_PVM_MASK) {
1731         if (msi_ctrl & PCI_MSI_64BIT_MASK)
1732             msi_mask_off = cap_ptr + PCI_MSI_64BIT_MASKBITS;
1733         else
1734             msi_mask_off = cap_ptr + PCI_MSI_32BIT_MASK;
1735         msi_pvm = pci_config_get32(handle, msi_mask_off);
1736         pci_config_put32(handle, msi_mask_off, (uint32_t)-1);
1737         APIC_VERBOSE(INTR, (CE_CONT,
1738             "set_grp: pvm supported. Mask set to 0x%x\n",
1739             pci_config_get32(handle, msi_mask_off)));
1740     }
1741
1742     if ((newp = apix_rebind(vecp, new_cpu, num_vectors)) != NULL)
1743         *result = 0;
1744
1745     /* Reenable vectors if per vector masking is supported. */
1746     if (msi_ctrl & PCI_MSI_PVM_MASK) {
1747         pci_config_put32(handle, msi_mask_off, msi_pvm);
1748         APIC_VERBOSE(INTR, (CE_CONT,
1749             "set_grp: pvm supported. Mask restored to 0x%x\n",
1750             pci_config_get32(handle, msi_mask_off)));
1751     }
1752
1753 }
1754 unchanged_portion_omitted_
```

```
new/usr/src/uts/i86pc/io/mp_platform_common.c
```

```
*****
```

```
68957 Thu Sep 7 15:25:32 2017
```

```
new/usr/src/uts/i86pc/io/mp_platform_common.c
```

```
8626 make pcplusmp and apix warning-free
```

```
Reviewed by: Robert Mustacchi <rm@joyent.com>
```

```
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
```

```
*****
```

```
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 2017 Joyent, Inc.
26 */
27 /*
28 * Copyright (c) 2010, Intel Corporation.
29 * All rights reserved.
30 */
31 /*
32 * PSMI 1.1 extensions are supported only in 2.6 and later versions.
33 * PSMI 1.2 extensions are supported only in 2.7 and later versions.
34 * PSMI 1.3 and 1.4 extensions are supported in Solaris 10.
35 * PSMI 1.5 extensions are supported in Solaris Nevada.
36 * PSMI 1.6 extensions are supported in Solaris Nevada.
37 * PSMI 1.7 extensions are supported in Solaris Nevada.
38 */
39 */
40 #define PSMI_1_7
41
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/sunddi.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>
```

```
1
```

```
new/usr/src/uts/i86pc/io/mp_platform_common.c
```

```
*****
```

```
60 #include <sys/uadmin.h>
61 #include <sys/panic.h>
62 #include <sys/debug.h>
63 #include <sys/archsystm.h>
64 #include <sys/trap.h>
65 #include <sys/machsystm.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/sundai.h>
73 #if !defined(_xpv)
74 #include <sys/hpet.h>
75 #include <sys/clock.h>
76#endif
77 /*
78 * Local Function Prototypes
79 */
80 static int apic_handle_defconf();
81 static int apic_parse_mpct(caddr_t mpct, int bypass);
82 static struct apic_mpfps_hdr *apic_find_fps_sig(caddr_t fptra, int size);
83 static int apic_checksum(caddr_t bptr, int len);
84 static int apic_find_bus_type(char *bus);
85 static int apic_find_bus(int busid);
86 static struct apic_io_intr *apic_find_io_intr(int irqno);
87 static int apic_find_free_irq(int start, int end);
88 static struct apic_io_intr *apic_find_io_intr_w_busid(int irqno, int busid);
89 static void apic_set_pwoff_method_from_mpcnfhdr(struct apic_mp_cnf_hdr *hdrrp);
90 static void apic_free_apic_cpus(void);
91 static boolean_t apic_is_iopanic_AMD_813x(uint32_t physaddr);
92 static int apic_acpi_enter_apicmode(void);
93
94 int apic_handle_pci_pci_bridge(dev_info_t *idip, int child_devno,
95 int child_ipin, struct apic_io_intr **intrp);
96 int apic_find_bus_id(int bustype);
97 int apic_find_intin(uchar_t ioapic, uchar_t intin);
98 void apic_record_rdt_entry(apic_irq_t *irqptr, int irq);
99
100 int apic_debug_mps_id = 0; /* 1 - print MPS ID strings */
101 /* ACPI SCI interrupt configuration; -1 if SCI not used */
102 int apic_sci_vect = -1;
103 iflag_t apic_sci_flags;
104
105 #if !defined(_xpv)
106 /* ACPI HPET interrupt configuration; -1 if HPET not used */
107 int apic_hpet_vect = -1;
108 iflag_t apic_hpet_flags;
109#endif
110
111 /*
112 * psm name pointer
113 */
114 char *psm_name;
115
116 /*
117 * ACPI support routines */
118 static int acpi_probe(char *);
119 static int acpi_configure(acpi_psm_lnk_t *acpipsmLnkP, dev_info_t *dip,
120 int *pci_irqp, iflag_t *intr_flagp);
121
122 int apic_acpi_translate_pci_irq(dev_info_t *dip, int busid, int devid,
123 int ipin, int *pci_irqp, iflag_t *intr_flagp);
124 uchar_t acpi_find_iopanic(int irq);
```

```
2
```

```

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 millisec */
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 *apicioaddr[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_inrp = 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_nexsts 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 * irqt 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_introp_xlate and programming in addspl.
257 */

```

```

258 kmutex_t      irq_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

311 int      apic_poweroff_method = APIC_POWEROFF_NONE;

313 /*
314  * Auto-configuration routines
315  */
317 /*
318  * Look at MPSpec 1.4 (Intel Order # 242016-005) for details of what we do here
319  * May work with 1.1 - but not guaranteed.
320  * According to the MP Spec, the MP floating pointer structure
321  * will be searched in the order described below:
322  * 1. In the first kilobyte of Extended BIOS Data Area (EBDA)
323  * 2. Within the last kilobyte of system base memory
324  * 3. In the BIOS ROM address space between 0F0000h and 0FFFFh
325  * Once we find the right signature with proper checksum, we call
326  * either handle_defconf or parse_mpct to get all info necessary for
327  * subsequent operations.
328 */
329 int
330 apic_probe_common(char *modname)
331 {
332     uint32_t mpct_addr, ebda_start = 0, base_mem_end;
333     caddr_t biosdatap;
334     caddr_t mpct = NULL;
335     caddr_t mpct = 0;
336     caddr_t fptr;
337     int i, mpct_size = 0, mapsize, retval = PSM_FAILURE;
338     ushort_t ebda_seg, base_mem_size;
339     struct apic_mpfps_hdr *fpfsp;
340     struct apic_mp_cnf_hdr *hdrp;
341     int bypass_cpu_and_ioapics_in_mptables;
342     int acpi_user_options;

343     if (apic_forceload < 0)
344         return (retval);

346     /*
347      * Remember who we are
348      */
349     psm_name = modname;

351     /* Allow override for MADT-only mode */
352     acpi_user_options = ddi_prop_get_int(DDI_DEV_T_ANY, ddi_root_node(), 0,
353                                         "acpi-user-options", 0);
354     apic_use_acpi_madt_only = ((acpi_user_options & ACPI_OUSER_MADT) != 0);

356     /* Allow apic_use_acpi to override MADT-only mode */
357     if (!apic_use_acpi)
358         apic_use_acpi_madt_only = 0;

360     retval = apci_probe(modname);

```

```

362     /* in UEFI system, there is no BIOS data */
363     if (ddi_prop_exists(DDI_DEV_T_ANY, ddi_root_node(), 0, "efi-systab"))
364         goto apic_ret;

366     /*
367      * mapin the bios data area 40:0
368      * 40:13h - two-byte location reports the base memory size
369      * 40:0Eh - two-byte location for the exact starting address of
370      *          the EBDA segment for EISA
371      */
372     biosdatap = psm_map_phys(0x400, 0x20, PROT_READ);
373     if (!biosdatap)
374         goto apic_ret;
375     fpfsp = (struct apic_mpfps_hdr *)NULL;
376     mapsize = MPFPS_RAM_WIN_LEN;
377     /*LINTED: pointer cast may result in improper alignment */
378     ebda_seg = *((ushort_t *) (biosdatap + 0xe));
379     /* check the 1k of EBDA */
380     if (ebda_seg) {
381         ebda_start = ((uint32_t)ebda_seg) << 4;
382         fptr = psm_map_phys(ebda_start, MPFPS_RAM_WIN_LEN, PROT_READ);
383         if (fptr) {
384             if (!(fpfsp =
385                   apic_find_fps_sig(fptr, MPFPS_RAM_WIN_LEN)))
386                 psm_unmap_phys(fptr, MPFPS_RAM_WIN_LEN);
387         }
388     }
389     /* If not in EBDA, check the last k of system base memory */
390     if (!fpfsp) {
391         /*LINTED: pointer cast may result in improper alignment */
392         base_mem_size = *((ushort_t *) (biosdatap + 0x13));

394         if (base_mem_size > 512)
395             base_mem_end = 639 * 1024;
396         else
397             base_mem_end = 511 * 1024;
398         /* if ebda == last k of base mem, skip to check BIOS ROM */
399         if (base_mem_end != ebda_start) {
400             fptr = psm_map_phys(base_mem_end, MPFPS_RAM_WIN_LEN,
401                                 PROT_READ);
402             if (fptr) {
403                 if (!(fpfsp =
404                       apic_find_fps_sig(fptr, MPFPS_RAM_WIN_LEN)))
405                     psm_unmap_phys(fptr, MPFPS_RAM_WIN_LEN);
406             }
407         }
408     }
409     psm_unmap_phys(biosdatap, 0x20);

413     /* If still cannot find it, check the BIOS ROM space */
414     if (!fpfsp) {
415         mapsize = MPFPS_ROM_WIN_LEN;
416         fptr = psm_map_phys(MPFPS_ROM_WIN_START,
417                             MPFPS_ROM_WIN_LEN, PROT_READ);
418         if (fptr) {
419             if (!(fpfsp =
420                   apic_find_fps_sig(fptr, MPFPS_ROM_WIN_LEN))) {
421                 psm_unmap_phys(fptr, MPFPS_ROM_WIN_LEN);
422             }
423         }
424     }
425 }
```

```

427     if (apic_checksum((caddr_t)fpfsp, fpfsp->mpfps_length * 16) != 0) {
428         psm_unmap_phys(fpstr, MPFPS_ROM_WIN_LEN);
429         goto apic_ret;
430     }
431
432     apic_spec_rev = fpfsp->mpfps_spec_rev;
433     if ((apic_spec_rev != 04) && (apic_spec_rev != 01)) {
434         psm_unmap_phys(fpstr, MPFPS_ROM_WIN_LEN);
435         goto apic_ret;
436     }
437
438     /* check IMCR is present or not */
439     apic_imcrp = fpfsp->mpfps_featinfo2 & MPFPS_FEATINFO2_IMCRP;
440
441     /* check default configuration (dual CPUs) */
442     if ((apic_defconf = fpfsp->mpfps_featinfo1) != 0) {
443         psm_unmap_phys(fpstr, mapsize);
444         if ((retval = apic_handle_defconf()) != PSM_SUCCESS)
445             return (retval);
446
447         goto apic_ret;
448     }
449
450     /* MP Configuration Table */
451     mpct_addr = (uint32_t)(fpfsp->mpfps_mpct_paddr);
452
453     psm_unmap_phys(fpstr, mapsize); /* unmap floating ptr struct */
454
455     /*
456      * Map in enough memory for the MP Configuration Table Header.
457      * Use this table to read the total length of the BIOS data and
458      * map in all the info
459      */
460     /*LINTED: pointer cast may result in improper alignment */
461     hdrp = (struct apic_mp_cnf_hdr *)psm_map_phys(mpct_addr,
462         sizeof (struct apic_mp_cnf_hdr), PROT_READ);
463     if (!hdrp)
464         goto apic_ret;
465
466     /* check mp configuration table signature PCMP */
467     if (hdrp->mpcnf_sig != 0x504d4350) {
468         psm_unmap_phys((caddr_t)hdrp, sizeof (struct apic_mp_cnf_hdr));
469         goto apic_ret;
470     }
471     mpct_size = (int)hdrp->mpcnf_tbl_length;
472
473     apic_set_pwroff_method_from_mpcnfhdr(hdrp);
474
475     psm_unmap_phys((caddr_t)hdrp, sizeof (struct apic_mp_cnf_hdr));
476
477     if ((retval == PSM_SUCCESS) && !apic_use_acpi_madt_only) {
478         /* This is an ACPI machine No need for further checks */
479         goto apic_ret;
480     }
481
482     /*
483      * Map in the entries for this machine, ie. Processor
484      * Entry Tables, Bus Entry Tables, etc.
485      * They are in fixed order following one another
486      */
487     mpct = psm_map_phys(mpct_addr, mpct_size, PROT_READ);
488     if (!mpct)
489         goto apic_ret;
490
491     if (apic_checksum(mpct, mpct_size) != 0)
492         goto apic_fail1;

```

```

494     /*LINTED: pointer cast may result in improper alignment */
495     hdrp = (struct apic_mp_cnf_hdr *)mpct;
496     apicadr = (uint32_t *)mapin_apic((uint32_t)hdrp->mpcnf_local_apic,
497         APIC_LOCAL_MEMLEN, PROT_READ | PROT_WRITE);
498     if (!apicadr)
499         goto apic_fail1;
500
501     /* Parse all information in the tables */
502     bypass_cpu_and_ioapics_in_mptables = (retval == PSM_SUCCESS);
503     if (apic_parse_mpct(mpct, bypass_cpu_and_ioapics_in_mptables) ==
504         PSM_SUCCESS) {
505         retval = PSM_SUCCESS;
506         goto apic_ret;
507     }
508
509 apic_fail1:
510     psm_unmap_phys(mpct, mpct_size);
511     mpct = NULL;
512
513 apic_ret:
514     if (retval == PSM_SUCCESS) {
515         extern int apic_ioapic_method_probe();
516
517         if ((retval = apic_ioapic_method_probe()) == PSM_SUCCESS)
518             return (PSM_SUCCESS);
519     }
520
521     for (i = 0; i < apic_io_max; i++)
522         mapout_ioapic((caddr_t)apicioaddr[i], APIC_IO_MEMLEN);
523     if (apic_cpus) {
524         kmem_free(apic_cpus, apic_cpus_size);
525         apic_cpus = NULL;
526     }
527     if (apicadr) {
528         mapout_apic((caddr_t)apicadr, APIC_LOCAL_MEMLEN);
529         apicadr = NULL;
530     }
531     if (mpct)
532         psm_unmap_phys(mpct, mpct_size);
533 }
534
535 } unchanged portion omitted

```

new/usr/src/uts/i86pc/io/mp_platform_misc.c

6396 Thu Sep 7 15:25:32 2017

new/usr/src/uts/i86pc/io/mp_platform_misc.c

8626 make pcplusmp and apix warning-free

Reviewed by: Robert Mustacchi <rm@joyent.com>

Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>

```
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) 2010, Oracle and/or its affiliates. All rights reserved.
23 * Copyright 2017 Joyent, Inc.
24 */
25 /*
26 * Copyright (c) 2010, Intel Corporation.
27 * All rights reserved.
28 */
29 /*
30 * PSMI 1.1 extensions are supported only in 2.6 and later versions.
31 * PSMI 1.2 extensions are supported only in 2.7 and later versions.
32 * PSMI 1.3 and 1.4 extensions are supported in Solaris 10.
33 * PSMI 1.5 extensions are supported in Solaris Nevada.
34 * PSMI 1.6 extensions are supported in Solaris Nevada.
35 * PSMI 1.7 extensions are supported in Solaris Nevada.
36 */
37 */
38 #define PSMI_1_7
39
40 #include <sys/processor.h>
41 #include <sys/time.h>
42 #include <sys/psm.h>
43 #include <sys/smp_impldefs.h>
44 #include <sys/inttypes.h>
45 #include <sys/cram.h>
46 #include <sys/acpi/acpi.h>
47 #include <sys/acpica.h>
48 #include <sys/psm_common.h>
49 #include <sys/apic.h>
50 #include <sys/apic_common.h>
51 #include <sys/pit.h>
52 #include <sys/ddi.h>
53 #include <sys/sunddi.h>
54 #include <sys/ddi_impldefs.h>
55 #include <sys/pci.h>
56 #include <sys/promif.h>
57 #include <sys/x86_archext.h>
58 #include <sys/epc_impl.h>
59 #include <sys/uadmin.h>
```

1

new/usr/src/uts/i86pc/io/mp_platform_misc.c

```
60 #include <sys/panic.h>
61 #include <sys/debug.h>
62 #include <sys/archsysm.h>
63 #include <sys/trap.h>
64 #include <sys/machsysm.h>
65 #include <sys/cpuvar.h>
66 #include <sys/rm_platter.h>
67 #include <sys/privregs.h>
68 #include <sys/cyclic.h>
69 #include <sys/note.h>
70 #include <sys/pci_intr_lib.h>
71 #include <sys/sunddi.h>
72 #include <sys/hpet.h>
73 #include <sys/clock.h>
74
75 /*
76 * Part of mp_platform_common.c that's used only by pcplusmp & xpv_psm
77 * but not apix.
78 * These functions may be moved to xpv_psm later when apix and pcplusmp
79 * are merged together
80 */
81
82 /*
83 * Local Function Prototypes
84 */
85 static void apic_mark_vector(uchar_t oldvector, uchar_t newvector);
86 static void apic_xlate_vector_free_timeout_handler(void *arg);
87 static int apic_check_stuck_interrupt(apic_irq_t *irq_ptr, int old_bind_cpu,
88     int new_bind_cpu, int apicindex, int intin_no, int which_irq,
89     struct ioapic_reprogram_data *drep);
90 static int apic_setup_irq_table(dev_info_t *dip, int irqno,
91     struct apic_io_intr *intrp, struct intrspec *ispec, iflag_t *intr_flagp,
92     int type);
93 static void apic_try_deferred_reprogram(int ipl, int vect);
94 static void delete_defer_repro_ent(int which_irq);
95 static void apic_ioapic_wait_pending_clear(int ioapicindex,
96     int intin_no);
97
98 extern int apic_acpi_translate_pci_irq(dev_info_t *dip, int busid, int devid,
99     int ipin, int *pci_irqp, iflag_t *intr_flagp);
100 extern int apic_handle_pci_pci_bridge(dev_info_t *idip, int child_devno,
101     int child_ipin, struct apic_io_intr **intrp);
102 extern uchar_t apic_find_ioapic(int irq);
103 extern struct apic_io_intr *apic_find_io_intr_w_busid(int irqno, int busid);
104 extern int apic_find_bus_id(int bustype);
105 extern int apic_find_intin(uchar_t ioapic, uchar_t intin);
106 extern void apic_record_rdt_entry(apic_irq_t *irqptr, int irq);
107
108 extern int apic_sci_vect;
109 extern iflag_t apic_sci_flags;
110 /* ACPI HPET interrupt configuration; -1 if HPET not used */
111 extern int apic_hpet_vect;
112 extern iflag_t apic_hpet_flags;
113 extern int apic_intr_policy;
114 extern char *psm_name;
115
116 /*
117 * number of bits per byte, from <sys/param.h>
118 */
119 #define UCHAR_MAX      UINT8_MAX
120
121 /* Max wait time (in repetitions) for flags to clear in an RDT entry. */
122 extern int apic_max_reps_clear_pending;
123
124 /* The irq # is implicit in the array index: */
125 struct ioapic_reprogram_data apic_reprogram_info[APIC_MAX_VECTOR+1];
```

2

```

126 /*
127  * APIC_MAX_VECTOR + 1 is the maximum # of IRQs as well. ioapic_reprogram_info
128  * is indexed by IRQ number, NOT by vector number.
129 */
130
131 extern int apic_int_busy_mark;
132 extern int apic_int_free_mark;
133 extern int apic_diff_for_redistribution;
134 extern int apic_sample_factor_redistribution;
135 extern int apic_redist_cpu_skip;
136 extern int apic_num_imbalance;
137 extern int apic_num_rebind;
138
139 /* timeout for xlate_vector, mark_vector */
140 int apic_revector_timeout = 16 * 10000; /* 160 millisec */
141
142 extern int apic_defconf;
143 extern int apic_irq_translate;
144
145 extern int apic_use_acpi_madt_only; /* 1=ONLY use MADT from ACPI */
146
147 extern uchar_t apic_io_vectbase[MAX_IO_APIC];
148
149 extern boolean_t ioapic_mask_workaround[MAX_IO_APIC];
150
151 /*
152  * First available slot to be used as IRQ index into the apic_irq_table
153  * for those interrupts (like MSI/X) that don't have a physical IRQ.
154 */
155 extern int apic_first_avail_irq;
156
157 /*
158  * apic_defer_reprogram_lock ensures that only one processor is handling
159  * deferred interrupt programming at *_intr_exit time.
160 */
161 static lock_t apic_defer_reprogram_lock;
162
163 /*
164  * The current number of deferred reprogrammings outstanding
165 */
166 uint_t apic_reprogram_outstanding = 0;
167
168 #ifdef DEBUG
169 /*
170  * Counters that keep track of deferred reprogramming stats
171 */
172 uint_t apic_intr_deferrals = 0;
173 uint_t apic_intr_deliver_timeouts = 0;
174 uint_t apic_last_ditch_reprogram_failures = 0;
175 uint_t apic_deferred_setup_failures = 0;
176 uint_t apic_defer_repro_total_retries = 0;
177 uint_t apic_defer_repro_successes = 0;
178 uint_t apic_deferred_spurious_enters = 0;
179#endif
180
181 extern int apic_io_max;
182 extern struct apic_io_intr *apic_io_intrp;
183
184 uchar_t apic_vector_to_irq[APIC_MAX_VECTOR+1];
185
186 extern uint32_t eisa_level_intr_mask;
187 /* At least MSB will be set if EISA bus */
188
189 extern int apic_pci_bus_total;
190 extern uchar_t apic_single_pci_busid;

```

```

192 /*
193  * Following declarations are for revectoring; used when ISRs at different
194  * IPLs share an irq.
195 */
196 static lock_t apic_revector_lock;
197 int apic_revector_pending = 0;
198 static uchar_t *apic_oldvec_to_newvec;
199 static uchar_t *apic_newvec_to_oldvec;
200
201 /* ACPI Interrupt Source Override Structure ptr */
202 extern ACPI_MADT_INTERRUPT_OVERRIDE *acpi_isop;
203 extern int acpi_iso_cnt;
204
205 /*
206  * Auto-configuration routines
207 */
208
209 /*
210  * Initialise vector->ipl and ipl->pri arrays. level_intr and irqtable
211  * are also set to NULL. vector->irq is set to a value which cannot map
212  * to a real irq to show that it is free.
213 */
214 void
215 apic_init_common(void)
216 {
217     int i, j, indx;
218     int *iptr;
219
220     /*
221      * Initialize apic_ipls from apic_vectortoipl. This array is
222      * used in apic_intr_enter to determine the IPL to use for the
223      * corresponding vector. On some systems, due to hardware errata
224      * and interrupt sharing, the IPL may not correspond to the IPL listed
225      * in apic_vectortoipl (see apic_addspl and apic_delspl).
226      */
227     for (i = 0; i < (APIC_AVAIL_VECTOR / APIC_VECTOR_PER_IPL); i++) {
228         indx = i * APIC_VECTOR_PER_IPL;
229
230         for (j = 0; j < APIC_VECTOR_PER_IPL; j++, indx++)
231             apic_ipls[indx] = apic_vectortoipl[i];
232     }
233
234     /* cpu 0 is always up (for now) */
235     apic_cpus[0].aci_status = APIC_CPU_ONLINE | APIC_CPU_INTR_ENABLE;
236
237     iptr = (int *)&apic_irq_table[0];
238     for (i = 0; i <= APIC_MAX_VECTOR; i++) {
239         apic_level_intr[i] = 0;
240         *iptr++ = NULL;
241         apic_vector_to_irq[i] = APIC_RESV_IRQ;
242
243         /* These *must* be inititted to B_TRUE! */
244         apic_reprogram_info[i].done = B_TRUE;
245         apic_reprogram_info[i].irqp = NULL;
246         apic_reprogram_info[i].tries = 0;
247         apic_reprogram_info[i].bindcpu = 0;
248     }
249
250     /*
251      * Allocate a dummy irq table entry for the reserved entry.
252      * This takes care of the race between removing an irq and
253      * clock detecting a CPU in that irq during interrupt load
254      * sampling.
255      */
256     apic_irq_table[APIC_RESV_IRQ] =
257         kmalloc(sizeof(apic_irq_t), KM_SLEEP);

```

```

259     mutex_init(&airq_mutex, NULL, MUTEX_DEFAULT, NULL);
260 }
unchanged_portion_omitted_
261
262 /* Recompute mask bits for the given interrupt vector.
263 * If there is no interrupt servicing routine for this
264 * vector, this function should disable interrupt vector
265 * from happening at all IPLs. If there are still
266 * handlers using the given vector, this function should
267 * disable the given vector from happening below the lowest
268 * IPL of the remaining hadlers.
269 */
270 /*ARGSUSED*/
271 int
272 apic_delspl_common(int irqno, int ipl, int min_ipl, int max_ipl)
273 {
274     uchar_t vector;
275     uint32_t bind_cpu;
276     int intin, irqindex;
277     int ioapic_ix;
278     apic_irq_t *irqptr, *preirqptr, *irqheadptr, *irqp;
279     ulong_t iflag;
280
281     mutex_enter(&airq_mutex);
282     irqindex = IRQINDEX(irqno);
283     irqptr = preirqptr = irqheadptr = apic_irq_table[irqindex];
284
285     DDI_INTR_IMPLDBG((CE_CONT, "apic_delspl: dip=0x%p type=%d irqno=0x%x "
286         "vector=0x%x\n", (void *)irqptr->airq_dip,
287         irqptr->airq_mps_intr_index, irqno, irqptr->airq_vector));
288
289     while (irqptr) {
290         if (VIRTIRQ(irqindex, irqptr->airq_share_id) == irqno)
291             break;
292         preirqptr = irqptr;
293         irqptr = irqptr->airq_next;
294     }
295     ASSERT(irqptr);
296
297     irqptr->airq_share--;
298
299     mutex_exit(&airq_mutex);
300
301     /*
302      * If there are more interrupts at a higher IPL, we don't need
303      * to disable anything.
304      */
305     if (ipl < max_ipl)
306         return (PSM_SUCCESS);
307
308     /* return if it is not hardware interrupt */
309     if (irqptr->airq_mps_intr_index == RESERVE_INDEX)
310         return (PSM_SUCCESS);
311
312     if (!apic_picinit_called) {
313         /*
314          * Clear irq_struct. If two devices shared an intpt
315          * line & l unloaded before picinit, we are hosed. But, then
316          * we hope the machine survive.
317          */
318         irqptr->airq_mps_intr_index = FREE_INDEX;
319         irqptr->airq_temp_cpu = IRQ_UNINIT;
320         apic_free_vector(irqptr->airq_vector);
321         return (PSM_SUCCESS);
322     }

```

```

323
324     }
325
326     /*
327      * Downgrade vector to new max_ipl if needed. If we cannot allocate,
328      * use old IPL. Not very elegant, but it should work.
329      */
330     if ((irqptr->airq_ipl != max_ipl) && (max_ipl != PSM_INVALID_IPL) &&
331         !ioapic_mask_workaround[irqptr->airq_ioapicindex]) {
332         apic_irq_t *irqp;
333         if ((vector = apic_allocate_vector(max_ipl, irqno, 1))) {
334             apic_mark_vector(irqheadptr->airq_vector, vector);
335             irqp = irqheadptr;
336             while (irqp) {
337                 irqp->airq_vector = vector;
338                 irqp->airq_ipl = (uchar_t)max_ipl;
339                 if (irqp->airq_temp_cpu != IRQ_UNINIT) {
340                     apic_record_rdt_entry(irqp, irqindex);
341
342                     iflag = intr_clear();
343                     lock_set(&apic_ioapic_lock);
344
345                     (void) apic_setup_io_intr(irqp,
346                         irqindex, B_FALSE);
347
348                     lock_clear(&apic_ioapic_lock);
349                     intr_restore(iflag);
350
351                     irqp = irqp->airq_next;
352                 }
353             }
354         } else if (irqptr->airq_ipl != max_ipl &&
355             max_ipl != PSM_INVALID_IPL &&
356             ioapic_mask_workaround[irqptr->airq_ioapicindex]) {
357
358         /*
359          * We cannot downgrade the IPL of the vector below the vector's
360          * hardware priority. If we did, it would be possible for a
361          * higher-priority hardware vector to interrupt a CPU running at an IPL
362          * lower than the hardware priority of the interrupting vector (but
363          * higher than the soft IPL of this IRQ). When this happens, we would
364          * then try to drop the IPL BELOW what it was (effectively dropping
365          * below base_spl) which would be potentially catastrophic.
366          *
367          * (e.g. Suppose the hardware vector associated with this IRQ is
368          * (hardware IPL of 4). Further assume that the old IPL of this IRQ
369          * was 4, but the new IPL is 1. If we forced vector 0x40 to result in
370          * an IPL of 1, it would be possible for the processor to be executing
371          * at IPL 3 and for an interrupt to come in on vector 0x40, interrupting
372          * the currently-executing ISR. When apic_intr_enter consults
373          * apic_irqs[], it will return 1, bringing the IPL of the CPU down to 1
374          * so even though the processor was running at IPL 4, an IPL 1
375          * interrupt will have interrupted it, which must not happen).
376          *
377          * Effectively, this means that the hardware priority corresponding to
378          * the IRQ's IPL (in apic_ipls[]) cannot be lower than the vector's
379          * hardware priority.
380          *
381          * (In the above example, then, after removal of the IPL 4 device's
382          * interrupt handler, the new IPL will continue to be 4 because the
383          * hardware priority that IPL 1 implies is lower than the hardware
384          * priority of the vector used.)
385          */
386         /* apic_ipls is indexed by vector, starting at APIC_BASE_VECT */
387         const int apic_ipls_index = irqptr->airq_vector -
388             APIC_BASE_VECT;
389         const int vect_inherent_hwpri = irqptr->airq_vector >>
390
391     }

```

```

615         APIC_IPL_SHIFT;
616
617     /*
618      * If there are still devices using this IRQ, determine the
619      * new ipl to use.
620     */
621     if (irqptr->airq_share) {
622         int vect_desired_hwpri, hwPRI;
623
624         ASSERT(max_ipl < MAXIPL);
625         vect_desired_hwpri = apic_ipltopri[max_ipl] >>
626             APIC_IPL_SHIFT;
627
628         /*
629          * If the desired IPL's hardware priority is lower
630          * than that of the vector, use the hardware priority
631          * of the vector to determine the new IPL.
632        */
633         hwPRI = (vect_desired_hwpri < vect_inherent_hwPRI) ?
634             vect_inherent_hwPRI : vect_desired_hwPRI;
635
636         /*
637          * Now, to get the right index for apic_vectortoipl,
638          * we need to subtract APIC_BASE_VECT from the
639          * hardware-vector-equivalent (in hwPRI). Since hwPRI
640          * is already shifted, we shift APIC_BASE_VECT before
641          * doing the subtraction.
642        */
643         hwPRI -= (APIC_BASE_VECT >> APIC_IPL_SHIFT);
644
645         ASSERT(hwPRI >= 0);
646         ASSERT(hwPRI < MAXIPL);
647         max_ipl = apic_vectortoipl[hwPRI];
648         apic_ipls[apic_ipls_index] = (uchar_t)max_ipl;
649         apic_ipls[apic_ipls_index] = max_ipl;
650
651         irqp = irqheadptr;
652         while (irqp) {
653             irqp->airq_ipl = (uchar_t)max_ipl;
654             irqp = irqp->airq_next;
655         } else {
656             /*
657              * No more devices on this IRQ, so reset this vector's
658              * element in apic_ipls to the original IPL for this
659              * vector
660            */
661             apic_ipls[apic_ipls_index] =
662                 apic_vectortoipl[vect_inherent_hwPRI];
663         }
664     }
665
666     /*
667      * If there are still active interrupts, we are done.
668     */
669     if (irqptr->airq_share)
670         return (PSM_SUCCESS);
671
672     iflag = intr_clear();
673     lock_set(&apic_ioapic_lock);
674
675     if (irqptr->airq_mps_intr_index == MSI_INDEX) {
676         /*
677          * Disable the MSI vector
678          * Make sure we only disable on the last
679          * of the multi-MSI support

```

```

680
681         /*
682          * If (i_ddi_intr_get_current_nenables(irqptr->airq_dip) == 1) {
683          *   apic_pci_msi_disable_mode(irqptr->airq_dip,
684          *                             DDI_INTR_TYPE_MSI);
685        }
686        } else if (irqptr->airq_mps_intr_index == MSIX_INDEX) {
687        /*
688          * Disable the MSI-X vector
689          * needs to clear its mask and addr/data for each MSI-X
690        */
691        apic_pci_msi_unconfigure(irqptr->airq_dip, DDI_INTR_TYPE_MSIX,
692                               irqptr->airq_origirq);
693        /*
694          * Make sure we only disable on the last MSI-X
695        */
696        if (i_ddi_intr_get_current_nenables(irqptr->airq_dip) == 1) {
697          apic_pci_msi_disable_mode(irqptr->airq_dip,
698                               DDI_INTR_TYPE_MSIX);
699        }
700      } else {
701        /*
702          * The assumption here is that this is safe, even for
703          * systems with IOAPICs that suffer from the hardware
704          * erratum because all devices have been quiesced before
705          * they unregister their interrupt handlers. If that
706          * assumption turns out to be false, this mask operation
707          * can induce the same erratum result we're trying to
708          * avoid.
709        */
710        ioapic_ix = irqptr->airq_ioapicindex;
711        intin = irqptr->airq_intin_no;
712        ioapic_write(ioapic_ix, APIC_RDT_CMD + 2 * intin, AV_MASK);
713    }
714
715    apic_vt_ops->apic_intrmap_free_entry(&irqptr->airq_intrmap_private);
716
717    /*
718      * This irq entry is the only one in the chain.
719    */
720    if (irqheadptr->airq_next == NULL) {
721        ASSERT(irqheadptr == irqptr);
722        bind_cpu = irqptr->airq_temp_cpu;
723        if (((uint32_t)bind_cpu != IRQ_UNBOUND) &&
724            ((uint32_t)bind_cpu != IRQ_UNINIT)) {
725            ASSERT(apic_cpu_in_range(bind_cpu));
726            if (bind_cpu & IRQ_USER_BOUND) {
727                /* If hardbound, temp_cpu == cpu */
728                bind_cpu &= ~IRQ_USER_BOUND;
729                apic_cpus[bind_cpu].aci_bound--;
730            } else
731                apic_cpus[bind_cpu].aci_temp_bound--;
732        }
733        irqptr->airq_temp_cpu = IRQ_UNINIT;
734        irqptr->airq_mps_intr_index = FREE_INDEX;
735        lock_clear(&apic_ioapic_lock);
736        intr_restore(iflag);
737        apic_free_vector(irqptr->airq_vector);
738        return (PSM_SUCCESS);
739    }
740
741    /*
742      * If we get here, we are sharing the vector and there are more than
743      * one active irq entries in the chain.
744    */
745    lock_clear(&apic_ioapic_lock);
746    intr_restore(iflag);

```

```

747     mutex_enter(&airq_mutex);
748     /* Remove the irq entry from the chain */
749     if (irqptr == irqheadptr) { /* The irq entry is at the head */
750         apic_irq_table[irqindex] = irqptr->airq_next;
751     } else {
752         preirqptr->airq_next = irqptr->airq_next;
753     }
754     /* Free the irq entry */
755     kmem_free(irqptr, sizeof (apic_irq_t));
756     mutex_exit(&airq_mutex);

758     return (PSM_SUCCESS);
759 }



---



unchanged_portion omitted



1042 /*
1043 * Allocate/Initialize the apic_irq_table[] entry for given irqno. If the entry
1044 * is used already, we will try to allocate a new irqno.
1045 *
1046 * Return value:
1047 *     Success: irqno
1048 *     Failure: -1
1049 */
1050 static int
1051 apic_setup_irq_table(dev_info_t *dip, int irqno, struct apic_io_intr *intrp,
1052                      struct intrspec *ispec, iflag_t *intr_flagp, int type)
1053 {
1054     int origirq;
1055     uchar_t ipl;
1056     int origirq = ispec->intrspec_vec;
1057     uchar_t ipin, ioapic, ioapicindex, vector;
1058     apic_irq_t *irqptr;
1059     major_t major;
1060     dev_info_t *sdip;

1062     ASSERT(ispec != NULL);

1064     origirq = ispec->intrspec_vec;
1065     ipl = ispec->intrspec_pri;

1067     DDI_INTR_IMPLDBG((CE_CONT, "apic_setup_irq_table: dip=0x%p type=%d "
1068                       "irqno=0x%x origirq=0x%x\n", (void *)dip, type, irqno, origirq));
1069
1070     ASSERT(ispec != NULL);

1072     if (DDI_INTR_IS_MSI_OR_MSIX(type)) {
1073         /* MSI/X doesn't need to setup ioapic stuffs */
1074         ioapicindex = 0xff;
1075         ioapic = 0xff;
1076         ipin = (uchar_t)0xff;
1077         intr_index = (type == DDI_INTR_TYPE_MSI) ? MSI_INDEX :
1078                         MSIX_INDEX;
1079         mutex_enter(&airq_mutex);
1080         if ((irqno = apic_allocate_irq(apic_first_avail_irq)) == -1) {
1081             mutex_exit(&airq_mutex);
1082             /* need an irq for MSI/X to index into autovect[] */
1083             cmn_err(CE_WARN, "No interrupt irq: %s instance %d",
1084                   ddi_get_name(dip), ddi_get_instance(dip));
1085             return (-1);
1086         }
1087         mutex_exit(&airq_mutex);
1088     }

```

```

1089     } else if (intrp != NULL) {
1090         intr_index = (int)(intrp - apic_io_intrp);
1091         ioapic = intrp->intr_destid;
1092         ipin = intrp->intr_destintin;
1093         /* Find ioapicindex. If destid was ALL, we will exit with 0. */
1094         for (ioapicindex = apic_io_max - 1; ioapicindex; ioapicindex--) {
1095             if (apic_io_id[ioapicindex] == ioapic)
1096                 break;
1097         }
1098         ASSERT((ioapic == apic_io_id[ioapicindex]) ||
1099                (ioapic == INTR_ALL_APIC));

1100        /* check whether this intin# has been used by another irqno */
1101        if ((newirq = apic_find_intin(ioapicindex, ipin)) != -1) {
1102            return (newirq);
1103        }

1105        } else if (intr_flagp != NULL) {
1106            /* ACPI case */
1107            intr_index = ACPI_INDEX;
1108            ioapicindex = acpi_find_ioapic(irqno);
1109            ASSERT(ioapicindex != 0xFF);
1110            ioapic = apic_io_id[ioapicindex];
1111            ipin = irqno - apic_io_vectbase[ioapicindex];
1112            if (apic_irq_table[irqno] &&
1113                apic_irq_table[irqno]->irq_mps_intr_index == ACPI_INDEX) {
1114                ASSERT(apic_irq_table[irqno]->irq_intin_no == ipin &
1115                      apic_irq_table[irqno]->irq_ioapicindex ==
1116                      ioapicindex);
1117                return (irqno);
1118            }
1120        } else {
1121            /* default configuration */
1122            ioapicindex = 0;
1123            ioapic = apic_io_id[ioapicindex];
1124            ipin = (uchar_t)irqno;
1125            intr_index = DEFAULT_INDEX;
1126        }

1128        if ((vector = apic_allocate_vector(ipl, irqno, 0)) == 0) {
1129            if (ispec == NULL) {
1130                APIC_VERBOSE_IOAPIC((CE_WARN, "No intrspec for irqno = %x\n",
1131                                     irqno));
1132            } else if ((vector = apic_allocate_vector(ipl, irqno, 0)) == 0) {
1133                if ((newirq = apic_share_vector(irqno, intr_flagp, intr_index,
1134                                                ipl, ioapicindex, ipin, &irqptr)) != -1) {
1135                    irqptr->airq_ipl = ipl;
1136                    irqptr->airq_origirq = (uchar_t)origirq;
1137                    irqptr->airq_ipin = ipin;
1138                    irqptr->airq_major = major;
1139                    sdip = apic_irq_table[IRQINDEX(newirq)]->airq_dip;
1140                    /* This is OK to do really */
1141                    if (sdip == NULL) {
1142                        cmn_err(CE_WARN, "Sharing vectors: %s"
1143                                " instance %d and SCI",
1144                                ddi_get_name(dip), ddi_get_instance(dip));
1145                } else {
1146                    cmn_err(CE_WARN, "Sharing vectors: %s"
1147                            " instance %d and %s instance %d",
1148                            ddi_get_name(sdip), ddi_get_instance(sdip),
1149                            ddi_get_name(dip), ddi_get_instance(dip));
1150                }
1151            }
1152            return (newirq);
1153        }
1154    }
1155    /* try high priority allocation now that share has failed */
1156

```

```

1150     if ((vector = apic_allocate_vector(ipl, irqno, 1)) == 0) {
1151         cmn_err(CE_WARN, "No interrupt vector: %s instance %d",
1152                 ddi_get_name(dip), ddi_get_instance(dip));
1153         return (-1);
1154     }
1155 }
1156 mutex_enter(&airq_mutex);
1157 if (apic_irq_table[irqno] == NULL) {
1158     irqptr = kmem_zalloc(sizeof (apic_irq_t), KM_SLEEP);
1159     irqptr->airq_temp_cpu = IRQ_UNINIT;
1160     apic_irq_table[irqno] = irqptr;
1161 } else {
1162     irqptr = apic_irq_table[irqno];
1163     if (irqptr->airq_mps_intr_index != FREE_INDEX) {
1164         /*
1165          * The slot is used by another irqno, so allocate
1166          * a free irqno for this interrupt
1167          */
1168         newirq = apic_allocate_irq(apic_first_avail_irq);
1169         if (newirq == -1) {
1170             mutex_exit(&airq_mutex);
1171             return (-1);
1172         }
1173         irqno = newirq;
1174         irqptr = apic_irq_table[irqno];
1175         if (irqptr == NULL) {
1176             irqptr = kmem_zalloc(sizeof (apic_irq_t),
1177                                 KM_SLEEP);
1178             irqptr->airq_temp_cpu = IRQ_UNINIT;
1179             apic_irq_table[irqno] = irqptr;
1180         }
1181     }
1182     vector = apic_modify_vector(vector, newirq);
1183 }
1184 apic_max_device_irq = max(irqno, apic_max_device_irq);
1185 apic_min_device_irq = min(irqno, apic_min_device_irq);
1186 mutex_exit(&airq_mutex);
1187 irqptr->airq_ioapicindex = ioapicindex;
1188 irqptr->airq_intin_no = ipin;
1189 irqptr->airq_ipl = ipl;
1190 irqptr->airq_vector = vector;
1191 irqptr->airq_origirq = (uchar_t)origirq;
1192 irqptr->airq_share_id = 0;
1193 irqptr->airq_mps_intr_index = (short)intr_index;
1194 irqptr->airq_dip = dip;
1195 irqptr->airq_major = major;
1196 irqptr->airq_cpu = apic_bind_intr(dip, irqno, ioapic, ipin);
1197 if (intr_flag)
1198     irqptr->airq_iflag = *intr_flag;
1199 if (!DDI_INTR_IS_MSI_OR_MSIX(type)) {
1200     /* setup I/O APIC entry for non-MSI/X interrupts */
1201     apic_record_rdt_entry(irqptr, irqno);
1202 }
1203 return (irqno);
1204 }
1205 */
1206 /*
1207 * return the cpu to which this intr should be bound.
1208 * Check properties or any other mechanism to see if user wants it
1209 * bound to a specific CPU. If so, return the cpu id with high bit set.
1210 * If not, use the policy to choose a cpu and return the id.
1211 */
1212 uint32_t
1213 apic_bind_intr(dev_info_t *dip, int irq, uchar_t ioapicid, uchar_t intin)

```

```

1216 {
1217     int instance, instno, prop_len, bind_cpu, count;
1218     uint_t i, rc;
1219     uint32_t cpu;
1220     major_t major;
1221     char *name, *drv_name, *prop_val, *cptr;
1222     char prop_name[32];
1223     ulong_t iflag;
1224
1225     if (apic_intr_policy == INTR_LOWEST_PRIORITY)
1226         return (IRQ_UNBOUND);
1227
1228     if (apic_nproc == 1)
1229         return (0);
1230
1231     if (dip == NULL) {
1232         iflag = intr_clear();
1233         lock_set(&apic_ioapic_lock);
1234         bind_cpu = apic_get_next_bind_cpu();
1235         lock_clear(&apic_ioapic_lock);
1236         intr_restore(iflag);
1237
1238         cmn_err(CE_CONT, "%s: irq 0x%lx "
1239                 "vector 0x%lx ioapic 0x%lx intin 0x%lx is bound to cpu %d\n",
1240                 psm_name, irq, apic_irq_table[irq]->airq_vector, ioapicid,
1241                 intin, bind_cpu & ~IRQ_USER_BOUND);
1242
1243         return ((uint32_t)bind_cpu);
1244     }
1245
1246     drv_name = NULL;
1247     rc = DDI_PROP_NOT_FOUND;
1248     major = (major_t)-1;
1249     if (dip != NULL) {
1250         name = ddi_get_name(dip);
1251         major = ddi_name_to_major(name);
1252         drv_name = ddi_major_to_name(major);
1253         instance = ddi_get_instance(dip);
1254         if (apic_intr_policy == INTR_ROUND_ROBIN_WITH_AFFINITY) {
1255             i = apic_min_device_irq;
1256             for (; i <= apic_max_device_irq; i++) {
1257                 if ((i == irq) || (apic_irq_table[i] == NULL) ||
1258                     (apic_irq_table[i]->airq_mps_intr_index
1259                      == FREE_INDEX))
1260                     continue;
1261
1262                 if ((apic_irq_table[i]->airq_major == major) &&
1263                     (!(apic_irq_table[i]->airq_cpu & IRQ_USER_BOUND)) &&
1264                     (!!(apic_irq_table[i]->airq_cpu &
1265                         IRQ_USER_BOUND))) {
1266                     cpu = apic_irq_table[i]->airq_cpu;
1267
1268                     cmn_err(CE_CONT,
1269                             "%!s: %s (%s) instance #%d "
1270                             "irq 0x%lx vector 0x%lx ioapic 0x%lx "
1271                             "intin 0x%lx is bound to cpu %d\n",
1272                             psm_name,
1273                             name, drv_name, instance, irq,
1274                             apic_irq_table[irq]->airq_vector,
1275                             ioapicid, intin, cpu);
1276                     return (cpu);
1277                 }
1278             }
1279         }
1280     }
1281
1282     if (apic_nproc == 1)
1283         return (0);
1284
1285     iflag = intr_clear();
1286     lock_set(&apic_ioapic_lock);
1287     bind_cpu = apic_get_next_bind_cpu();
1288     lock_clear(&apic_ioapic_lock);
1289     intr_restore(iflag);
1290
1291     cmn_err(CE_CONT,
1292             "%!s: %s (%s) instance #%d "
1293             "irq 0x%lx vector 0x%lx ioapic 0x%lx "
1294             "intin 0x%lx is bound to cpu %d\n",
1295             psm_name,
1296             name, drv_name, instance, irq,
1297             apic_irq_table[irq]->airq_vector,
1298             ioapicid, intin, cpu);
1299     return (bind_cpu);
1300 }

```

```

1274     }
1275     /*
1276      * search for "drvname"_intpt_bind_cpus property first, the
1277      * syntax of the property should be "a[,b,c,...]" where
1278      * instance 0 binds to cpu a, instance 1 binds to cpu b,
1279      * instance 3 binds to cpu c...
1280      * ddi_getlongprop() will search /option first, then /
1281      * if "drvname"_intpt_bind_cpus doesn't exist, then find
1282      * intpt_bind_cpus property. The syntax is the same, and
1283      * it applies to all the devices if its "drvname" specific
1284      * property doesn't exist
1285     */
1286     (void) strcpy(prop_name, drv_name);
1287     (void) strcat(prop_name, "_intpt_bind_cpus");
1288     rc = ddi_getlongprop(DDI_DEV_T_ANY, dip, 0, prop_name,
1289                          (caddr_t)&prop_val, &prop_len);
1290     if (rc != DDI_PROP_SUCCESS) {
1291         rc = ddi_getlongprop(DDI_DEV_T_ANY, dip, 0,
1292                             "intpt_bind_cpus", (caddr_t)&prop_val, &prop_len);
1293     }
1294     if (rc == DDI_PROP_SUCCESS) {
1295         for (i = count = 0; i < (prop_len - 1); i++)
1296             if (prop_val[i] == ',')
1297                 count++;
1298         if (prop_val[i-1] != ',')
1299             count++;
1300         /*
1301          * if somehow the binding instances defined in the
1302          * property are not enough for this instno, then
1303          * reuse the pattern for the next instance until
1304          * it reaches the requested instno
1305         */
1306         instno = instance % count;
1307         i = 0;
1308         cptr = prop_val;
1309         while (i < instno)
1310             if (*cptr++ == ',')
1311                 i++;
1312         bind_cpu = atoi(&cptr);
1313         kmem_free(prop_val, prop_len);
1314         /* if specific CPU is bogus, then default to next cpu */
1315         if (!apic_cpu_in_range(bind_cpu)) {
1316             cmn_err(CE_WARN, "%s: %s=%s: CPU %d not present",
1317                     psm_name, prop_name, prop_val, bind_cpu);
1318             rc = DDI_PROP_NOT_FOUND;
1319         } else {
1320             /* indicate that we are bound at user request */
1321             bind_cpu |= IRQ_USER_BOUND;
1322         }
1323         /*
1324          * no need to check apic_cpus[].aci_status, if specific CPU is
1325          * not up, then post_cpu_start will handle it.
1326         */
1327     }
1328     if (rc != DDI_PROP_SUCCESS) {
1329         iflag = intr_clear();
1330         lock_set(&apic_ioapic_lock);
1331         bind_cpu = apic_get_next_bind_cpu();
1332         lock_clear(&apic_ioapic_lock);
1333         intr_restore(iflag);
1334     }
1335
1328     if (drv_name != NULL)
1336         cmn_err(CE_CONT, "%s (%s) instance %d irq 0x%x "

```

```

1338         "vector 0x%x ioapic 0x%x intin 0x%x is bound to cpu %d\n",
1339         psm_name, name, drv_name, instance, irq,
1340         apic_irq_table[irq]->irq_vector, ioapicid, intin,
1341         bind_cpu & ~IRQ_USER_BOUND);
1334     else
1335         cmn_err(CE_CONT, "%s: irq 0x%x "
1336         "vector 0x%x ioapic 0x%x intin 0x%x is bound to cpu %d\n",
1337         psm_name, irq, apic_irq_table[irq]->irq_vector, ioapicid,
1338         intin, bind_cpu & ~IRQ_USER_BOUND);
1343     return ((uint32_t)bind_cpu);
1344 }
```

unchanged portion omitted

new/usr/src/uts/i86pc/io/pcplusmp/apic.c

1

```
*****
34798 Thu Sep 7 15:25:32 2017
new/usr/src/uts/i86pc/io/pcplusmp/apic.c
8626 make pcplusmp and apix warning-free
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
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) 1993, 2010, Oracle and/or its affiliates. All rights reserved.
23  */
24 /*
25  * Copyright (c) 2010, Intel Corporation.
26  * All rights reserved.
27  */
28 /*
29  * Copyright (c) 2017, Joyent, Inc. All rights reserved.
30  */
31 /*
32  *
33  * To understand how the pcplusmp module interacts with the interrupt subsystem
34  * read the theory statement in uts/i86pc/os/intr.c.
35  */
36 /*
37  *
38  * PSMI 1.1 extensions are supported only in 2.6 and later versions.
39  * PSMI 1.2 extensions are supported only in 2.7 and later versions.
40  * PSMI 1.3 and 1.4 extensions are supported in Solaris 10.
41  * PSMI 1.5 extensions are supported in Solaris Nevada.
42  * PSMI 1.6 extensions are supported in Solaris Nevada.
43  * PSMI 1.7 extensions are supported in Solaris Nevada.
44  */
45 /*
46 #define PSMI_1_7
47
48 #include <sys/processor.h>
49 #include <sys/time.h>
50 #include <sys/psm.h>
51 #include <sys/smp_ImplDefs.h>
52 #include <sys/cram.h>
53 #include <sys/acpi/acpi.h>
54 #include <sys/acpica.h>
55 #include <sys/psm_common.h>
56 #include <sys/apic.h>
57 #include <sys/pit.h>
58 #include <sys/ddi.h>
59 #include <sys/sunddi.h>
```

new/usr/src/uts/i86pc/io/pcplusmp/apic.c

2

```
60 #include <sys/ddi_ImplDefs.h>
61 #include <sys/pci.h>
62 #include <sys/promif.h>
63 #include <sys/x86_archext.h>
64 #include <sys/cpc_impl.h>
65 #include <sys/uadmin.h>
66 #include <sys/panic.h>
67 #include <sys/debug.h>
68 #include <sys/archsystm.h>
69 #include <sys/trap.h>
70 #include <sys/machsystm.h>
71 #include <sys/sysmacros.h>
72 #include <sys/cpuvar.h>
73 #include <sys/rm_platter.h>
74 #include <sys/privregs.h>
75 #include <sys/note.h>
76 #include <sys/pci_intr_lib.h>
77 #include <sys/spl.h>
78 #include <sys/clock.h>
79 #include <sys/cyclic.h>
80 #include <sys/dditypes.h>
81 #include <sys/sunddi.h>
82 #include <sys/x_call.h>
83 #include <sys/reboot.h>
84 #include <sys/hpet.h>
85 #include <sys/apic_common.h>
86 #include <sys/apic_timer.h>
87
88 /*
89  * Local Function Prototypes
90  */
91 static void apic_init_intr(void);
92
93 /*
94  * standard MP entries
95  */
96 static int apic_probe(void);
97 static int apic_getclkirq(int ipl);
98 static void apic_init(void);
99 static void apic_picinit(void);
100 static int apic_post_cpu_start(void);
101 static int apic_intr_enter(int ipl, int *vect);
102 static void apic_setspl(int ipl);
103 static void x2apic_setspl(int ipl);
104 static int apic_addspl(int ipl, int vector, int min_ipl, int max_ipl);
105 static int apic_delspl(int ipl, int vector, int min_ipl, int max_ipl);
106 static int apic_disable_intr(processorid_t cpun);
107 static void apic_enable_intr(processorid_t cpun);
108 static int apic_get_ipivect(int ipl, int type);
109 static void apic_post_cyclic_setup(void *arg);
110
111 #define UCHAR_MAX     UINT8_MAX
112
113 /*
114  * The following vector assignments influence the value of ipltopri and
115  * vectortoipl. Note that vectors 0 - 0x1f are not used. We can program
116  * idle to 0 and IPL 0 to 0xf to differentiate idle in case
117  * we care to do so in future. Note some IPLs which are rarely used
118  * will share the vector ranges and heavily used IPLs (5 and 6) have
119  * a wide range.
120  *
121  * This array is used to initialize apic_ipls[] (in apic_init()).
122  *
123  *          IPL           Vector range.      as passed to intr_enter
124  *          0             none.
125  *          1,2,3         0x20-0x2f        0x0-0xf
```

```

126 *      4          0x30-0x3f      0x10-0x1f
127 *      5          0x40-0x5f      0x20-0x3f
128 *      6          0x60-0x7f      0x40-0x5f
129 *      7,8,9     0x80-0x8f      0x60-0x6f
130 *      10         0x90-0x9f      0x70-0x7f
131 *      11         0xa0-0xaf      0x80-0x8f
132 *      ...
133 *      15         0xe0-0xef      0xc0-0xcf
134 *      15         0xf0-0xff      0xd0-0xdf
135 */
136 uchar_t apic_vectortoipl[APIC_AVAIL_VECTOR / APIC_VECTOR_PER_IPL] = {
137     3, 4, 5, 5, 6, 6, 9, 10, 11, 12, 13, 14, 15, 15
138 };
unchanged portion omitted

1064 /*
1065 * This function allocates "count" MSI vector(s) for the given "dip/pri/type"
1066 */
1067 int
1068 apic_alloc_msi_vectors(dev_info_t *dip, int inum, int count, int pri,
1069                         int behavior)
1070 {
1071     int rcount, i;
1072     uchar_t start, irqno;
1073     uint32_t cpu = 0;
1074     uint32_t cpu;
1075     major_t major;
1076     apic_irq_t *irqptr;
1077
1078     DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msi_vectors: dip=0x%p "
1079                     "inum=0x%x pri=0x%x count=0x%x behavior=%d\n",
1080                     (void *)dip, inum, pri, count, behavior));
1081
1082     if (count > 1) {
1083         if (behavior == DDI_INTR_ALLOC_STRICT &&
1084             apic_multi_msi_enable == 0)
1085             return (0);
1086         if (apic_multi_msi_enable == 0)
1087             count = 1;
1088     }
1089
1090     if ((rcount = apic_navail_vector(dip, pri)) > count)
1091         rcount = count;
1092     else if (rcount == 0 || (rcount < count &&
1093                             behavior == DDI_INTR_ALLOC_STRICT))
1094         return (0);
1095
1096     /* if not ISP2, then round it down */
1097     if (!ISP2(rcount))
1098         rcount = 1 << (highbit(rcount) - 1);
1099
1100     mutex_enter(&airq_mutex);
1101
1102     for (start = 0; rcount > 0; rcount >>= 1) {
1103         if ((start = apic_find_multi_vectors(pri, rcount)) != 0 ||
1104             behavior == DDI_INTR_ALLOC_STRICT)
1105             break;
1106     }
1107
1108     if (start == 0) {
1109         /* no vector available */
1110         mutex_exit(&airq_mutex);
1111         return (0);
1112     }
1113
1114     if (apic_check_free_irqs(rcount) == PSM_FAILURE) {

```

```

1114             /* not enough free irq slots available */
1115             mutex_exit(&airq_mutex);
1116             return (0);
1117         }
1118
1119         major = (dip != NULL) ? ddi_driver_major(dip) : 0;
1120         for (i = 0; i < rcount; i++) {
1121             if ((irqno = apic_allocate_irq(apic_first_avail_irq)) ==
1122                 (uchar_t)-1) {
1123                 /*
1124                  * shouldn't happen because of the
1125                  * apic_check_free_irqs() check earlier
1126                  */
1127                 mutex_exit(&airq_mutex);
1128                 DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msi_vectors: "
1129                                 "apic_allocate_irq failed\n"));
1130                 return (i);
1131             }
1132             apic_max_device_irq = max(irqno, apic_max_device_irq);
1133             apic_min_device_irq = min(irqno, apic_min_device_irq);
1134             irqptr = apic_irq_table[irqno];
1135 #ifdef DEBUG
1136             if (apic_vector_to_irq[start + i] != APIC_RESV IRQ)
1137                 DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msi_vectors: "
1138                                 "apic_vector_to_irq is not APIC_RESV IRQ\n"));
1139 #endif
1140             apic_vector_to_irq[start + i] = (uchar_t)irqno;
1141
1142             irqptr->airq_vector = (uchar_t)(start + i);
1143             irqptr->airq_ioapicindex = (uchar_t)inum; /* start */
1144             irqptr->airq_intin_no = (uchar_t)rcount;
1145             ASSERT(pri >= 0 && pri <= UCHAR_MAX);
1146             irqptr->airq_irlp = (uchar_t)pri;
1147             irqptr->airq_iprl = pri;
1148             irqptr->airq_vector = start + i;
1149             irqptr->airq_origirq = (uchar_t)(inum + i);
1150             irqptr->airq_share_id = 0;
1151             irqptr->airq_mps_intr_index = MSI_INDEX;
1152             irqptr->airq_dip = dip;
1153             irqptr->airq_major = major;
1154             if (i == 0) /* they all bound to the same cpu */
1155                 cpu = irqptr->airq_cpu = apic_bind_intr(dip, irqno,
1156                                                     0xff, 0xff);
1157             else
1158                 irqptr->airq_cpu = cpu;
1159             DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msi_vectors: irq=0x%x "
1160                             "dip=0x%p vector=0x%x origirq=0x%x pri=0x%x\n",
1161                             (void *)irqptr->airq_dip, irqptr->airq_vector,
1162                             irqptr->airq_origirq, pri));
1163             mutex_exit(&airq_mutex);
1164         }
1165
1166     /* This function allocates "count" MSI-X vector(s) for the given "dip/pri/type"
1167 */
1168     int
1169     apic_alloc_msix_vectors(dev_info_t *dip, int inum, int count, int pri,
1170                             int behavior)
1171 {
1172     int rcount, i;
1173     major_t major;
1174
1175     mutex_enter(&airq_mutex);

```

```

1179     if ((rcount = apic_navail_vector(dip, pri)) > count)
1180         rcount = count;
1181     else if (rcount == 0 || (rcount < count &&
1182         behavior == DDI_INTR_ALLOC_STRICT)) {
1183         rcount = 0;
1184         goto out;
1185     }
1186
1187     if (apic_check_free_irqs(rcount) == PSM_FAILURE) {
1188         /* not enough free irq slots available */
1189         rcount = 0;
1190         goto out;
1191     }
1192
1193     major = (dip != NULL) ? ddi_driver_major(dip) : 0;
1194     for (i = 0; i < rcount; i++) {
1195         uchar_t vector, irqno;
1196         apic_irq_t *irqptr;
1197
1198         if ((irqno = apic_allocate_irq(apic_first_avail_irq)) ==
1199             (uchar_t)-1) {
1200             /*
1201             * shouldn't happen because of the
1202             * apic_check_free_irqs() check earlier
1203             */
1204             DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msix_vectors: "
1205                             "apic_allocate_irq failed\n"));
1206             rcount = i;
1207             goto out;
1208         }
1209         if ((vector = apic_allocate_vector(pri, irqno, 1)) == 0) {
1210             /*
1211             * shouldn't happen because of the
1212             * apic_navail_vector() call earlier
1213             */
1214             DDI_INTR_IMPLDBG((CE_CONT, "apic_alloc_msix_vectors: "
1215                             "apic_allocate_vector failed\n"));
1216             rcount = i;
1217             goto out;
1218         }
1219         apic_max_device_irq = max(irqno, apic_max_device_irq);
1220         apic_min_device_irq = min(irqno, apic_min_device_irq);
1221         irqptr = apic_irq_table[irqno];
1222         irqptr->airq_vector = (uchar_t)vector;
1223         ASSERT(pri > 0 && pri <= UCHAR_MAX);
1224         irqptr->airq_ipl = (uchar_t)pri;
1225         irqptr->airq_ipl = pri;
1226         irqptr->airq_origirq = (uchar_t)(inum + i);
1227         irqptr->airq_share_id = 0;
1228         irqptr->airq_mps_intr_index = MSIX_INDEX;
1229         irqptr->airq_dip = dip;
1230         irqptr->airq_major = major;
1231         irqptr->airq_cpu = apic_bind_intr(dip, irqno, 0xff, 0xff);
1232     }
1233 out:
1234     mutex_exit(&airq_mutex);
1235     return (rcount);
1236 }
1237 */
1238 * Allocate a free vector for irq at ipl. Takes care of merging of multiple
1239 * IPLs into a single APIC level as well as stretching some IPLs onto multiple
1240 * levels. APIC_HI_PRI_VECTS interrupts are reserved for high priority
1241 * requests and allocated only when pri is set.
1242 */
1243 uchar_t

```

```

1244 apic_allocate_vector(int ipl, int irq, int pri)
1245 {
1246     int      lowest, highest, i;
1247
1248     highest = apic_ipltopri[ipl] + APIC_VECTOR_MASK;
1249     lowest = apic_ipltopri[ipl - 1] + APIC_VECTOR_PER_IPL;
1250
1251     if (highest < lowest) /* Both ipl and ipl - 1 map to same pri */
1252         lowest -= APIC_VECTOR_PER_IPL;
1253
1254 #ifdef DEBUG
1255     if (apic_restrict_vector) /* for testing shared interrupt logic */
1256         highest = lowest + apic_restrict_vector + APIC_HI_PRI_VECTS;
1257 #endif /* DEBUG */
1258     if (pri == 0)
1259         highest -= APIC_HI_PRI_VECTS;
1260
1261     for (i = lowest; i <= highest; i++) {
1262         if (APIC_CHECK_RESERVE_VECTORS(i))
1263             continue;
1264         if (apic_vector_to_irq[i] == APIC_RESV_IRQ) {
1265             apic_vector_to_irq[i] = (uchar_t)irq;
1266             ASSERT(i >= 0 && i <= UCHAR_MAX);
1267             return ((uchar_t)i);
1268         }
1269     }
1270
1271     return (0);
1272 }
1273
1274 unchanged portion omitted

```

```
*****
4487 Thu Sep 7 15:25:32 2017
new/usr/src/uts/i86pc/io/pcplusmp/apic_common.c
8626 make pcplusmp and apix warning-free
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
unchanged_portion_omitted_
```

```
845 int
846 apic_cpu_add(psm_cpu_request_t *reqp)
847 {
848     int i, rv = 0;
849     ulong_t iflag;
850     boolean_t first = B_TRUE;
851     uchar_t localver = 0;
851     uchar_t localver;
852     uint32_t localid, procid;
853     processorid_t cpuid = (processorid_t)-1;
854     mach_cpu_add_arg_t *ap;
855
856     ASSERT(reqp != NULL);
857     reqp->req.cpu_add.cpuid = (processorid_t)-1;
858
859     /* Check whether CPU hotplug is supported. */
860     if (!plat_dr_support_cpu() || apic_max_nproc == -1) {
861         return (ENOTSUP);
862     }
863
864     ap = (mach_cpu_add_arg_t *)reqp->req.cpu_add.argp;
865     switch (ap->type) {
866     case MACH_CPU_ARG_LOCAL_APIC:
867         localid = ap->arg.apic.apic_id;
868         procid = ap->arg.apic.proc_id;
869         if (localid >= 255 || procid > 255) {
870             cmn_err(CE_WARN,
871                     "!apic: apicid(%u) or procid(%u) is invalid.", localid, procid);
872             return (EINVAL);
873         }
874         break;
875
876     case MACH_CPU_ARG_LOCAL_X2APIC:
877         localid = ap->arg.apic.apic_id;
878         procid = ap->arg.apic.proc_id;
879         if (localid >= UINT32_MAX) {
880             cmn_err(CE_WARN,
881                     "!apic: x2apicid(%u) is invalid.", localid);
882             return (EINVAL);
883         } else if (localid >= 255 && apic_mode == LOCAL_APIC) {
884             cmn_err(CE_WARN, "!apic: system is in APIC mode, "
885                     "can't support x2APIC processor.");
886             return (ENOTSUP);
887         }
888         break;
889
890     default:
891         cmn_err(CE_WARN,
892                 "!apic: unknown argument type %d to apic_cpu_add()", ap->type);
893         return (EINVAL);
894     }
895
896     /* Use apic_ioapic_lock to sync with apic_get_next_bind_cpu. */
897     iflag = intr_clear();
898     lock_set(&apic_ioapic_lock);
```

```
902     /* Check whether local APIC id already exists. */
903     for (i = 0; i < apic_nproc; i++) {
904         if (!CPU_IN_SET(apic_cpumask, i))
905             continue;
906         if (apic_cpus[i].aci_local_id == localid) {
907             lock_clear(&apic_ioapic_lock);
908             intr_restore(iflag);
909             cmn_err(CE_WARN,
910                     "!apic: local apic id %u already exists.", localid);
911             return (EEXIST);
912         } else if (apic_cpus[i].aci_processor_id == procid) {
913             lock_clear(&apic_ioapic_lock);
914             intr_restore(iflag);
915             cmn_err(CE_WARN,
916                     "!apic: processor id %u already exists.", (int)procid);
917             return (EEXIST);
918         }
919     }
920
921     /*
922      * There's no local APIC version number available in MADT table,
923      * so assume that all CPUs are homogeneous and use local APIC
924      * version number of the first existing CPU.
925      */
926     if (first) {
927         first = B_FALSE;
928         localver = apic_cpus[i].aci_local_ver;
929     }
930     ASSERT(first == B_FALSE);
931
932
933     /*
934      * Try to assign the same cpuid if APIC id exists in the dirty cache.
935      */
936     for (i = 0; i < apic_max_nproc; i++) {
937         if (CPU_IN_SET(apic_cpumask, i)) {
938             ASSERT((apic_cpus[i].aci_status & APIC_CPU_FREE) == 0);
939             continue;
940         }
941         ASSERT(apic_cpus[i].aci_status & APIC_CPU_FREE);
942         if ((apic_cpus[i].aci_status & APIC_CPU_DIRTY) &&
943             apic_cpus[i].aci_local_id == localid &&
944             apic_cpus[i].aci_processor_id == procid) {
945             cpuid = i;
946             break;
947         }
948     }
949
950     /*
951      * Avoid the dirty cache and allocate fresh slot if possible.
952      */
953     if (cpuid == (processorid_t)-1) {
954         for (i = 0; i < apic_max_nproc; i++) {
955             if ((apic_cpus[i].aci_status & APIC_CPU_FREE) &&
956                 (apic_cpus[i].aci_status & APIC_CPU_DIRTY) == 0) {
957                 cpuid = i;
958                 break;
959             }
960         }
961
962     /*
963      * Try to find any free slot as last resort.
964      */
965     if (cpuid == (processorid_t)-1) {
966         for (i = 0; i < apic_max_nproc; i++) {
967             if (apic_cpus[i].aci_status & APIC_CPU_FREE) {
968                 cpuid = i;
```

```
967         }
968     }
969 }
970
972 if (cpuid == (processorid_t)-1) {
973     lock_clear(&apic_ioapic_lock);
974     intr_restore(iflag);
975     cmn_err(CE_NOTE,
976             "!apic: failed to allocate cpu id for processor %u.",
977             procid);
978     rv = EAGAIN;
979 } else if (ACPI_FAILURE(acpica_map_cpu(cpuid, procid))) {
980     lock_clear(&apic_ioapic_lock);
981     intr_restore(iflag);
982     cmn_err(CE_NOTE,
983             "!apic: failed to build mapping for processor %u.",
984             procid);
985     rv = EBUSY;
986 } else {
987     ASSERT(cpuid >= 0 && cpuid < NCPU);
988     ASSERT(cpuid < apic_max_nproc && cpuid < max_ncpus);
989     bzero(&apic_cpus[cpuid], sizeof (apic_cpus[0]));
990     apic_cpus[cpuid].aci_processor_id = procid;
991     apic_cpus[cpuid].aci_local_id = localid;
992     apic_cpus[cpuid].aci_local_ver = localver;
993     CPUSET_ATOMIC_ADD(apic_cpumask, cpuid);
994     if (cpuid >= apic_nproc) {
995         apic_nproc = cpuid + 1;
996     }
997     lock_clear(&apic_ioapic_lock);
998     intr_restore(iflag);
999     reqp->req.cpu_add.cpuid = cpuid;
1000 }
1002
1003 }
```

unchanged portion omitted

new/usr/src/uts/i86pc/io/pcplusmp/apic_introp.c

```
*****
26149 Thu Sep 7 15:25:32 2017
new/usr/src/uts/i86pc/io/pcplusmp/apic_introp.c
8626 make pcplusmp and apix warning-free
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
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) 2005, 2010, Oracle and/or its affiliates. All rights reserved.
23 * Copyright 2013 Pluribus Networks, Inc.
24 * Copyright 2017 Joyent, Inc.
25 */

26 /*
27 * apic_introp.c:
28 *      Has code for Advanced DDI interrupt framework support.
29 */
30 */

31 #include <sys/cpuvar.h>
32 #include <sys/psm.h>
33 #include <sys/archsystm.h>
34 #include <sys/apic.h>
35 #include <sys/sunddi.h>
36 #include <sys/ddi_impldefs.h>
37 #include <sys/mach_intr.h>
38 #include <sys/sysmacros.h>
39 #include <sys/trap.h>
40 #include <sys/pci.h>
41 #include <sys/pci_intr_lib.h>
42 #include <sys/apic_common.h>

43 #define UCHAR_MAX     UINT8_MAX
44
45 extern struct av_head autovect[];
46
47 /*
48 * Local Function Prototypes
49 */
50 apic_irq_t *apic_find_irq(dev_info_t *, struct intrspec *, int);
51
52 apic_pci_msi_enable_vector:
53     Set the address/data fields in the MSI/X capability structure
54     XXX: MSI-X support
55 */
56
57 /* ARGSUSED */
58
```

1

new/usr/src/uts/i86pc/io/pcplusmp/apic_introp.c

```
*****
60 void
61 apic_pci_msi_enable_vector(apic_irq_t *irq_ptr, int type, int inum, int vector,
62                             int count, int target_apic_id)
63 {
64     uint64_t msi_addr, msi_data;
65     ushort_t msi_ctrl;
66     dev_info_t dev_info_t
67     int dip = irq_ptr->airq_dip;
68     ddi_acc_handle_t cap_ptr = i_ddi_get_msi_msix_cap_ptr(dip);
69     handle = i_ddi_get_pci_config_handle(dip);
70     msi_regs_t msi_regs;
71     int irqno, i;
72     void *intrmap_tbl[PCI_MSI_MAX_INTRS];

73     DDI_INTR_IMPLDBG((CE_CONT, "apic_pci_msi_enable_vector: dip=0x%p\n"
74                         "\tdriver = %s, inum=0x%x vector=0x%x apicid=0x%x\n", (void *)dip,
75                         ddi_driver_name(dip), inum, vector, target_apic_id));

76     ASSERT((handle != NULL) && (cap_ptr != 0));

77     msi_regs.mr_data = vector;
78     msi_regs.mr_addr = target_apic_id;

79     for (i = 0; i < count; i++) {
80         irqno = apic_vector_to_irq(vector + i);
81         intrmap_tbl[i] = apic_irq_table[irqno]->airq_intrmap_private;
82     }
83     apic_vt_ops->apic_intrmap_alloc_entry(intrmap_tbl, dip, type,
84                                         count, 0xff);
85     for (i = 0; i < count; i++) {
86         irqno = apic_vector_to_irq(vector + i);
87         apic_irq_table[irqno]->airq_intrmap_private =
88             intrmap_tbl[i];
89     }
90
91     apic_vt_ops->apic_intrmap_map_entry(irq_ptr->airq_intrmap_private,
92                                         (void *)&msi_regs, type, count);
93     apic_vt_ops->apic_intrmap_record_msi(irq_ptr->airq_intrmap_private,
94                                         &msi_regs);

95     /* MSI Address */
96     msi_addr = msi_regs.mr_addr;
97
98     /* MSI Data: MSI is edge triggered according to spec */
99     msi_data = msi_regs.mr_data;
100
101    DDI_INTR_IMPLDBG((CE_CONT, "apic_pci_msi_enable_vector: addr=0x%lx "
102                      "data=0x%lx\n", (long)msi_addr, (long)msi_data));
103
104    if (type == DDI_INTR_TYPE_MSI) {
105        msi_ctrl = pci_config_get16(handle, cap_ptr + PCI_MSI_CTRL);
106
107        /* Set the bits to inform how many MSIs are enabled */
108        msi_ctrl |= ((highbit(count) - 1) << PCI_MSI_MME_SHIFT);
109        pci_config_put16(handle, cap_ptr + PCI_MSI_CTRL, msi_ctrl);

110        /*
111         * Only set vector if not on hypervisor
112         */
113        pci_config_put32(handle,
114                         cap_ptr + PCI_MSI_ADDR_OFFSET, msi_addr);

115        if (msi_ctrl & PCI_MSI_64BIT_MASK) {
116            pci_config_put32(handle,
117                             cap_ptr + PCI_MSI_ADDR_OFFSET + 4, msi_addr >> 32);
118            pci_config_put16(handle,
119                             cap_ptr + PCI_MSI_64BIT_DATA, msi_data);
120        }
121    }
122
123
```

2

```

126         } else {
127             pci_config_put16(handle,
128                             cap_ptr + PCI_MSI_32BIT_DATA, msi_data);
129         }
130     } else if (type == DDI_INTR_TYPE_MSIX) {
131         uintptr_t off;
132         ddi_intr_msix_t *msix_p = i_ddi_get_msix(dip);
133
134         ASSERT(msix_p != NULL);
135
136         /* Offset into the "inum"th entry in the MSI-X table */
137         off = (uintptr_t)msix_p->msix_tbl_addr +
138               (inum * PCI_MSIX_VECTOR_SIZE);
139
140         ddi_put32(msix_p->msix_tbl_hdl,
141                   (uint32_t *)(off + PCI_MSIX_DATA_OFFSET), msi_data);
142         ddi_put32(msix_p->msix_tbl_hdl,
143                   (uint32_t *)(off + PCI_MSIX_LOWER_ADDR_OFFSET), msi_addr);
144         ddi_put32(msix_p->msix_tbl_hdl,
145                   (uint32_t *)(off + PCI_MSIX_UPPER_ADDR_OFFSET),
146                   msi_addr >> 32);
147     }
148 }
149
150 unchanged_portion_omitted
151
152 /* Finds "count" contiguous MSI vectors starting at the proper alignment
153  * at "pri".
154  * Caller needs to make sure that count has to be power of 2 and should not
155  * be < 1.
156  */
157 uchar_t
158 apic_find_multi_vectors(int pri, int count)
159 {
160     int      lowest, highest, i, navail, start, msibits;
161
162     DDI_INTR_IMPLDBG((CE_CONT, "apic_find_mult: pri: %x, count: %x\n",
163                      pri, count));
164
165     highest = apic_ipltopri[pri] + APIC_VECTOR_MASK;
166     lowest = apic_ipltopri[pri - 1] + APIC_VECTOR_PER_IPL;
167     navail = 0;
168
169     if (highest < lowest) /* Both ipl and ipl - 1 map to same pri */
170         lowest -= APIC_VECTOR_PER_IPL;
171
172     /*
173      * msibits is the no. of lower order message data bits for the
174      * allocated MSI vectors and is used to calculate the aligned
175      * starting vector
176      */
177     msibits = count - 1;
178
179     /* It has to be contiguous */
180     for (i = lowest; i <= highest; i++) {
181         navail = 0;
182
183         /*
184          * starting vector has to be aligned accordingly for
185          * multiple MSIs
186          */
187         if (msibits)
188             i = (i + msibits) & ~msibits;
189         start = i;
190         while ((apic_vector_to_irq[i] == APIC_RESV_IRQ) &&
191                (i <= highest)) {
192             pci_config_put16(handle,
193                             cap_ptr + PCI_MSI_32BIT_DATA, msi_data);
194
195             ASSERT(msix_p != NULL);
196
197             /* Offset into the "inum"th entry in the MSI-X table */
198             off = (uintptr_t)msix_p->msix_tbl_addr +
199                   (inum * PCI_MSIX_VECTOR_SIZE);
200
201             ddi_put32(msix_p->msix_tbl_hdl,
202                       (uint32_t *)(off + PCI_MSIX_DATA_OFFSET), msi_data);
203             ddi_put32(msix_p->msix_tbl_hdl,
204                       (uint32_t *)(off + PCI_MSIX_LOWER_ADDR_OFFSET), msi_addr);
205             ddi_put32(msix_p->msix_tbl_hdl,
206                       (uint32_t *)(off + PCI_MSIX_UPPER_ADDR_OFFSET),
207                       msi_addr >> 32);
208         }
209     }
210 }
211
212 unchanged_portion_omitted
213
214 /* Finds "count" contiguous MSI vectors starting at the proper alignment
215  * at "pri".
216  * Caller needs to make sure that count has to be power of 2 and should not
217  * be < 1.
218  */
219 uchar_t
220 apic_find_multi_vectors(int pri, int count)
221 {
222     int      lowest, highest, i, navail, start, msibits;
223
224     highest = apic_ipltopri[pri] + APIC_VECTOR_MASK;
225     lowest = apic_ipltopri[pri - 1] + APIC_VECTOR_PER_IPL;
226     navail = 0;
227
228     if (highest < lowest) /* Both ipl and ipl - 1 map to same pri */
229         lowest -= APIC_VECTOR_PER_IPL;
230
231     /*
232      * msibits is the no. of lower order message data bits for the
233      * allocated MSI vectors and is used to calculate the aligned
234      * starting vector
235      */
236     msibits = count - 1;
237
238     /* It has to be contiguous */
239     for (i = lowest; i <= highest; i++) {
240         navail = 0;
241
242         /*
243          * starting vector has to be aligned accordingly for
244          * multiple MSIs
245          */
246         if (msibits)
247             i = (i + msibits) & ~msibits;
248         start = i;
249         while ((apic_vector_to_irq[i] == APIC_RESV_IRQ) &&
250                (i <= highest)) {
251             pci_config_put16(handle,
252                             cap_ptr + PCI_MSI_32BIT_DATA, msi_data);
253
254             ASSERT(msix_p != NULL);
255
256             /* Offset into the "inum"th entry in the MSI-X table */
257             off = (uintptr_t)msix_p->msix_tbl_addr +
258                   (inum * PCI_MSIX_VECTOR_SIZE);
259
260             ddi_put32(msix_p->msix_tbl_hdl,
261                       (uint32_t *)(off + PCI_MSIX_DATA_OFFSET), msi_data);
262             ddi_put32(msix_p->msix_tbl_hdl,
263                       (uint32_t *)(off + PCI_MSIX_LOWER_ADDR_OFFSET), msi_addr);
264             ddi_put32(msix_p->msix_tbl_hdl,
265                       (uint32_t *)(off + PCI_MSIX_UPPER_ADDR_OFFSET),
266                       msi_addr >> 32);
267         }
268     }
269 
```

```

270         (i <= highest)) {
271             if (APIC_CHECK_RESERVE_VECTORS(i))
272                 break;
273             navail++;
274         }
275         if (navail >= count) {
276             ASSERT(start >= 0 && start <= UCHAR_MAX);
277             return ((uchar_t)start);
278         }
279         if (navail >= count)
280             return (start);
281         i++;
282     }
283 }
284
285 unchanged_portion_omitted
286
287 static int
288 apic_grp_set_cpu(int irqno, int new_cpu, int *result)
289 {
290     dev_info_t *orig_dip;
291     uint32_t orig_cpu;
292     ulong_t iflag;
293     apic_irq_t *irqps[PCI_MSI_MAX_INTRS];
294     int i;
295     int cap_ptr;
296
297     int msi_mask_off = 0;
298     int msi_mask_off;
299     ushort_t msi_ctrl;
300     uint32_t msi_pvm = 0;
301     uint32_t msi_pvm;
302     ddi_acc_handle_t handle;
303     int num_vectors = 0;
304     uint32_t vector;
305
306     DDI_INTR_IMPLDBG((CE_CONT, "APIC_GRP_SET_CPU\n"));
307
308     /*
309      * Take mutex to insure that table doesn't change out from underneath
310      * us while we're playing with it.
311      */
312     mutex_enter(&airq_mutex);
313     irqps[0] = apic_irq_table[irqno];
314     orig_cpu = irqps[0]->airq_temp_cpu;
315     orig_dip = irqps[0]->airq_dip;
316     num_vectors = irqps[0]->airq_intin_no;
317     vector = irqps[0]->airq_vector;
318
319     /* A "group" of 1 */
320     if (num_vectors == 1) {
321         mutex_exit(&airq_mutex);
322         return (apic_set_cpu(irqno, new_cpu, result));
323     }
324
325     *result = ENXIO;
326
327     if (irqps[0]->airq_mps_intr_index != MSI_INDEX) {
328         mutex_exit(&airq_mutex);
329         DDI_INTR_IMPLDBG((CE_CONT, "set_grp: intr not MSI\n"));
330         goto set_grp_intr_done;
331     }
332     if ((num_vectors < 1) || ((num_vectors - 1) & vector)) {
333         mutex_exit(&airq_mutex);
334         DDI_INTR_IMPLDBG((CE_CONT,
335                           "set_grp: base vec not part of a grp or not aligned: "
336                           "vec:0x%x, num_vec:0x%x\n", vector, num_vectors));
337     }
338
339     /* Set up interrupt table */
340     if (num_vectors > 1) {
341         apic_set_group(irqps[0], new_cpu);
342         apic_set_group(irqps[0], orig_dip);
343         apic_set_group(irqps[0], orig_cpu);
344         apic_set_group(irqps[0], num_vectors);
345         apic_set_group(irqps[0], vector);
346     }
347
348     /* Set up interrupt table */
349     if (num_vectors > 1) {
350         apic_set_group(irqps[0], new_cpu);
351         apic_set_group(irqps[0], orig_dip);
352         apic_set_group(irqps[0], orig_cpu);
353         apic_set_group(irqps[0], num_vectors);
354         apic_set_group(irqps[0], vector);
355     }
356
357     /* Set up interrupt table */
358     if (num_vectors > 1) {
359         apic_set_group(irqps[0], new_cpu);
360         apic_set_group(irqps[0], orig_dip);
361         apic_set_group(irqps[0], orig_cpu);
362         apic_set_group(irqps[0], num_vectors);
363         apic_set_group(irqps[0], vector);
364     }
365
366     /* Set up interrupt table */
367     if (num_vectors > 1) {
368         apic_set_group(irqps[0], new_cpu);
369         apic_set_group(irqps[0], orig_dip);
370         apic_set_group(irqps[0], orig_cpu);
371         apic_set_group(irqps[0], num_vectors);
372         apic_set_group(irqps[0], vector);
373     }
374
375     /* Set up interrupt table */
376     if (num_vectors > 1) {
377         apic_set_group(irqps[0], new_cpu);
378         apic_set_group(irqps[0], orig_dip);
379         apic_set_group(irqps[0], orig_cpu);
380         apic_set_group(irqps[0], num_vectors);
381         apic_set_group(irqps[0], vector);
382     }
383
384     /* Set up interrupt table */
385     if (num_vectors > 1) {
386         apic_set_group(irqps[0], new_cpu);
387         apic_set_group(irqps[0], orig_dip);
388         apic_set_group(irqps[0], orig_cpu);
389         apic_set_group(irqps[0], num_vectors);
390         apic_set_group(irqps[0], vector);
391     }
392
393     /* Set up interrupt table */
394     if (num_vectors > 1) {
395         apic_set_group(irqps[0], new_cpu);
396         apic_set_group(irqps[0], orig_dip);
397         apic_set_group(irqps[0], orig_cpu);
398         apic_set_group(irqps[0], num_vectors);
399         apic_set_group(irqps[0], vector);
400     }
401
402     /* Set up interrupt table */
403     if (num_vectors > 1) {
404         apic_set_group(irqps[0], new_cpu);
405         apic_set_group(irqps[0], orig_dip);
406         apic_set_group(irqps[0], orig_cpu);
407         apic_set_group(irqps[0], num_vectors);
408         apic_set_group(irqps[0], vector);
409     }
410
411     /* Set up interrupt table */
412     if (num_vectors > 1) {
413         apic_set_group(irqps[0], new_cpu);
414         apic_set_group(irqps[0], orig_dip);
415         apic_set_group(irqps[0], orig_cpu);
416         apic_set_group(irqps[0], num_vectors);
417         apic_set_group(irqps[0], vector);
418     }
419
420     /* Set up interrupt table */
421     if (num_vectors > 1) {
422         apic_set_group(irqps[0], new_cpu);
423         apic_set_group(irqps[0], orig_dip);
424         apic_set_group(irqps[0], orig_cpu);
425         apic_set_group(irqps[0], num_vectors);
426         apic_set_group(irqps[0], vector);
427     }
428
429     /* Set up interrupt table */
430     if (num_vectors > 1) {
431         apic_set_group(irqps[0], new_cpu);
432         apic_set_group(irqps[0], orig_dip);
433         apic_set_group(irqps[0], orig_cpu);
434         apic_set_group(irqps[0], num_vectors);
435         apic_set_group(irqps[0], vector);
436     }
437
438     /* Set up interrupt table */
439     if (num_vectors > 1) {
440         apic_set_group(irqps[0], new_cpu);
441         apic_set_group(irqps[0], orig_dip);
442         apic_set_group(irqps[0], orig_cpu);
443         apic_set_group(irqps[0], num_vectors);
444         apic_set_group(irqps[0], vector);
445     }
446
447     /* Set up interrupt table */
448     if (num_vectors > 1) {
449         apic_set_group(irqps[0], new_cpu);
450         apic_set_group(irqps[0], orig_dip);
451         apic_set_group(irqps[0], orig_cpu);
452         apic_set_group(irqps[0], num_vectors);
453         apic_set_group(irqps[0], vector);
454     }
455
456     /* Set up interrupt table */
457     if (num_vectors > 1) {
458         apic_set_group(irqps[0], new_cpu);
459         apic_set_group(irqps[0], orig_dip);
460         apic_set_group(irqps[0], orig_cpu);
461         apic_set_group(irqps[0], num_vectors);
462         apic_set_group(irqps[0], vector);
463     }
464
465     /* Set up interrupt table */
466     if (num_vectors > 1) {
467         apic_set_group(irqps[0], new_cpu);
468         apic_set_group(irqps[0], orig_dip);
469         apic_set_group(irqps[0], orig_cpu);
470         apic_set_group(irqps[0], num_vectors);
471         apic_set_group(irqps[0], vector);
472     }
473
474     /* Set up interrupt table */
475     if (num_vectors > 1) {
476         apic_set_group(irqps[0], new_cpu);
477         apic_set_group(irqps[0], orig_dip);
478         apic_set_group(irqps[0], orig_cpu);
479         apic_set_group(irqps[0], num_vectors);
480         apic_set_group(irqps[0], vector);
481     }
482
483     /* Set up interrupt table */
484     if (num_vectors > 1) {
485         apic_set_group(irqps[0], new_cpu);
486         apic_set_group(irqps[0], orig_dip);
487         apic_set_group(irqps[0], orig_cpu);
488         apic_set_group(irqps[0], num_vectors);
489         apic_set_group(irqps[0], vector);
490     }
491
492     /* Set up interrupt table */
493     if (num_vectors > 1) {
494         apic_set_group(irqps[0], new_cpu);
495         apic_set_group(irqps[0], orig_dip);
496         apic_set_group(irqps[0], orig_cpu);
497         apic_set_group(irqps[0], num_vectors);
498         apic_set_group(irqps[0], vector);
499     }
500
501     /* Set up interrupt table */
502     if (num_vectors > 1) {
503         apic_set_group(irqps[0], new_cpu);
504         apic_set_group(irqps[0], orig_dip);
505         apic_set_group(irqps[0], orig_cpu);
506         apic_set_group(irqps[0], num_vectors);
507         apic_set_group(irqps[0], vector);
508     }
509
510     /* Set up interrupt table */
511     if (num_vectors > 1) {
512         apic_set_group(irqps[0], new_cpu);
513         apic_set_group(irqps[0], orig_dip);
514         apic_set_group(irqps[0], orig_cpu);
515         apic_set_group(irqps[0], num_vectors);
516         apic_set_group(irqps[0], vector);
517     }
518
519     /* Set up interrupt table */
520     if (num_vectors > 1) {
521         apic_set_group(irqps[0], new_cpu);
522         apic_set_group(irqps[0], orig_dip);
523         apic_set_group(irqps[0], orig_cpu);
524         apic_set_group(irqps[0], num_vectors);
525         apic_set_group(irqps[0], vector);
526     }
527
528     /* Set up interrupt table */
529     if (num_vectors > 1) {
530         apic_set_group(irqps[0], new_cpu);
531         apic_set_group(irqps[0], orig_dip);
532         apic_set_group(irqps[0], orig_cpu);
533         apic_set_group(irqps[0], num_vectors);
534         apic_set_group(irqps[0], vector);
535     }
536
537     /* Set up interrupt table */
538     if (num_vectors > 1) {
539         apic_set_group(irqps[0], new_cpu);
540         apic_set_group(irqps[0], orig_dip);
541         apic_set_group(irqps[0], orig_cpu);
542         apic_set_group(irqps[0], num_vectors);
543         apic_set_group(irqps[0], vector);
544     }
545
546     /* Set up interrupt table */
547     if (num_vectors > 1) {
548         apic_set_group(irqps[0], new_cpu);
549         apic_set_group(irqps[0], orig_dip);
550         apic_set_group(irqps[0], orig_cpu);
551         apic_set_group(irqps[0], num_vectors);
552         apic_set_group(irqps[0], vector);
553     }
554
555     /* Set up interrupt table */
556     if (num_vectors > 1) {
557         apic_set_group(irqps[0], new_cpu);
558         apic_set_group(irqps[0], orig_dip);
559         apic_set_group(irqps[0], orig_cpu);
560         apic_set_group(irqps[0], num_vectors);
561         apic_set_group(irqps[0], vector);
562     }
563
564     /* Set up interrupt table */
565     if (num_vectors > 1) {
566         apic_set_group(irqps[0], new_cpu);
567         apic_set_group(irqps[0], orig_dip);
568         apic_set_group(irqps[0], orig_cpu);
569         apic_set_group(irqps[0], num_vectors);
570         apic_set_group(irqps[0], vector);
571     }
572
573     /* Set up interrupt table */
574     if (num_vectors > 1) {
575         apic_set_group(irqps[0], new_cpu);
576         apic_set_group(irqps[0], orig_dip);
577         apic_set_group(irqps[0], orig_cpu);
578         apic_set_group(irqps[0], num_vectors);
579         apic_set_group(irqps[0], vector);
580     }
581
582     /* Set up interrupt table */
583     if (num_vectors > 1) {
584         apic_set_group(irqps[0], new_cpu);
585         apic_set_group(irqps[0], orig_dip);
586         apic_set_group(irqps[0], orig_cpu);
587         apic_set_group(irqps[0], num_vectors);
588         apic_set_group(irqps[0], vector);
589     }
590
591     /* Set up interrupt table */
592     if (num_vectors > 1) {
593         apic_set_group(irqps[0], new_cpu);
594         apic_set_group(irqps[0], orig_dip);
595         apic_set_group(irqps[0], orig_cpu);
596         apic_set_group(irqps[0], num_vectors);
597         apic_set_group(irqps[0], vector);
598     }
599
600     /* Set up interrupt table */
601     if (num_vectors > 1) {
602         apic_set_group(irqps[0], new_cpu);
603         apic_set_group(irqps[0], orig_dip);
604         apic_set_group(irqps[0], orig_cpu);
605         apic_set_group(irqps[0], num_vectors);
606         apic_set_group(irqps[0], vector);
607     }
608
609     /* Set up interrupt table */
610     if (num_vectors > 1) {
611         apic_set_group(irqps[0], new_cpu);
612         apic_set_group(irqps[0], orig_dip);
613         apic_set_group(irqps[0], orig_cpu);
614         apic_set_group(irqps[0], num_vectors);
615         apic_set_group(irqps[0], vector);
616     }
617
618     /* Set up interrupt table */
619     if (num_vectors > 1) {
620         apic_set_group(irqps[0], new_cpu);
621         apic_set_group(irqps[0], orig_dip);
622         apic_set_group(irqps[0], orig_cpu);
623         apic_set_group(irqps[0], num_vectors);
624         apic_set_group(irqps[0], vector);
625     }
626
627     /* Set up interrupt table */
628     if (num_vectors > 1) {
629         apic_set_group(irqps[0], new_cpu);
630         apic_set_group(irqps[0], orig_dip);
631         apic_set_group(irqps[0], orig_cpu);
632         apic_set_group(irqps[0], num_vectors);
633         apic_set_group(irqps[0], vector);
634     }
635
636     /* Set up interrupt table */
637     if (num_vectors > 1) {
638         apic_set_group(irqps[0], new_cpu);
639         apic_set_group(irqps[0], orig_dip);
640         apic_set_group(irqps[0], orig_cpu);
641         apic_set_group(irqps[0], num_vectors);
642         apic_set_group(irqps[0], vector);
643     }
644
645     /* Set up interrupt table */
646     if (num_vectors > 1) {
647         apic_set_group(irqps[0], new_cpu);
648         apic_set_group(irqps[0], orig_dip);
649         apic_set_group(irqps[0], orig_cpu);
650         apic_set_group(irqps[0], num_vectors);
651         apic_set_group(irqps[0], vector);
652     }
653
654     /* Set up interrupt table */
655     if (num_vectors > 1) {
656         apic_set_group(irqps[0], new_cpu);
657         apic_set_group(irqps[0], orig_dip);
658         apic_set_group(irqps[0], orig_cpu);
659         apic_set_group(irqps[0], num_vectors);
660         apic_set_group(irqps[0], vector);
661     }
662
663     /* Set up interrupt table */
664     if (num_vectors > 1) {
665         apic_set_group(irqps[0], new_cpu);
666         apic_set_group(irqps[0], orig_dip);
667         apic_set_group(irqps[0], orig_cpu);
668         apic_set_group(irqps[0], num_vectors);
669         apic_set_group(irqps[0], vector);
670     }
671
672     /* Set up interrupt table */
673     if (num_vectors > 1) {
674         apic_set_group(irqps[0], new_cpu);
675         apic_set_group(irqps[0], orig_dip);
676         apic_set_group(irqps[0], orig_cpu);
677         apic_set_group(irqps[0], num_vectors);
678         apic_set_group(irqps[0], vector);
679     }
680
681     /* Set up interrupt table */
682     if (num_vectors > 1) {
683         apic_set_group(irqps[0], new_cpu);
684         apic_set_group(irqps[0], orig_dip);
685         apic_set_group(irqps[0], orig_cpu);
686         apic_set_group(irqps[0], num_vectors);
687         apic_set_group(irqps[0], vector);
688     }
689
690     /* Set up interrupt table */
691     if (num_vectors > 1) {
692         apic_set_group(irqps[0], new_cpu);
693         apic_set_group(irqps[0], orig_dip);
694         apic_set_group(irqps[0], orig_cpu);
695         apic_set_group(irqps[0], num_vectors);
696         apic_set_group(irqps[0], vector);
697     }
698
699     /* Set up interrupt table */
700     if (num_vectors > 1) {
701         apic_set_group(irqps[0], new_cpu);
702         apic_set_group(irqps[0], orig_dip);
703         apic_set_group(irqps[0], orig_cpu);
704         apic_set_group(irqps[0], num_vectors);
705         apic_set_group(irqps[0], vector);
706     }
707
708     /* Set up interrupt table */
709     if (num_vectors > 1) {
710         apic_set_group(irqps[0], new_cpu);
711         apic_set_group(irqps[0], orig_dip);
712         apic_set_group(irqps[0], orig_cpu);
713         apic_set_group(irqps[0], num_vectors);
714         apic_set_group(irqps[0], vector);
715     }
716
717     /* Set up interrupt table */
718     if (num_vectors > 1) {
719         apic_set_group(irqps[0], new_cpu);
720         apic_set_group(irqps[0], orig_dip);
721         apic_set_group(irqps[0], orig_cpu);
722         apic_set_group(irqps[0], num_vectors);
723         apic_set_group(irqps[0], vector);
724     }
725
726     /* Set up interrupt table */
727     if (num_vectors > 1) {
728         apic_set_group(irqps[0], new_cpu);
729         apic_set_group(irqps[0], orig_dip);
730         apic_set_group(irqps[0], orig_cpu);
731         apic_set_group(irqps[0], num_vectors);
732         apic_set_group(irqps[0], vector);
733     }
734
735     /* Set up interrupt table */
736     if (num_vectors > 1) {
737         apic_set_group(irqps[0], new_cpu);
738         apic_set_group(irqps[0], orig_dip);
739         apic_set_group(irqps[0], orig_cpu);
740         apic_set_group(irqps[0], num_vectors);
741         apic_set_group(irqps[0], vector);
742     }
743
744     /* Set up interrupt table */
745     if (num_vectors > 1) {
746         apic_set_group(irqps[0], new_cpu);
747         apic_set_group(irqps[0], orig_dip);
748         apic_set_group(irqps[0], orig_cpu);
749         apic_set_group(irqps[0], num_vectors);
750         apic_set_group(irqps[0], vector);
751     }
752
753     /* Set up interrupt table */
754     if (num_vectors > 1) {
755         apic_set_group(irqps[0], new_cpu);
756         apic_set_group(irqps[0], orig_dip);
757         apic_set_group(irqps[0], orig_cpu);
758         apic_set_group(irqps[0], num_vectors);
759         apic_set_group(irqps[0], vector);
760     }
761
762     /* Set up interrupt table */
763     if (num_vectors > 1) {
764         apic_set_group(irqps[0], new_cpu);
765         apic_set_group(irqps[0], orig_dip);
766         apic_set_group(irqps[0], orig_cpu);
767         apic_set_group(irqps[0], num_vectors);
768         apic_set_group(irqps[0], vector);
769     }
770
771     /* Set up interrupt table */
772     if (num_vectors > 1) {
773         apic_set_group(irqps[0], new_cpu);
774         apic_set_group(irqps[0], orig_dip);
775         apic_set_group(irqps[0], orig_cpu);
776         apic_set_group(irqps[0], num_vectors);
777         apic_set_group(irqps[0], vector);
778     }
779
780     /* Set up interrupt table */
781     if (num_vectors > 1) {
782         apic_set_group(irqps[0], new_cpu);
783         apic_set_group(irqps[0], orig_dip);
784         apic_set_group(irqps[0], orig_cpu);
785         apic_set_group(irqps[0], num_vectors);
786         apic_set_group(irqps[0], vector);
787     }
788
789     /* Set up interrupt table */
790     if (num_vectors > 1) {
791         apic_set_group(irqps[0], new_cpu);
792         apic_set_group(irqps[0], orig_dip);
793         apic_set_group(irqps[0], orig_cpu);
794         apic_set_group(irqps[0], num_vectors);
795         apic_set_group(irqps[0], vector);
796     }
797
798     /* Set up interrupt table */
799     if (num_vectors > 1) {
800         apic_set_group(irqps[0], new_cpu);
801         apic_set_group(irqps[0], orig_dip);
802         apic_set_group(irqps[0], orig_cpu);
803         apic_set_group(irqps[0], num_vectors);
804         apic_set_group(irqps[0], vector);
805     }
806
807     /* Set up interrupt table */
808     if (num_vectors > 1) {
809         apic_set_group(irqps[0], new_cpu);
810         apic_set_group(irqps[0], orig_dip);
811         apic_set_group(irqps[0], orig_cpu);
812         apic_set_group(irqps[0], num_vectors);
813         apic_set_group(irqps[0], vector);
814     }
815
816     /* Set up interrupt table */
817     if (num_vectors > 1) {
818         apic_set_group(irqps[0], new_cpu);
819         apic_set_group(irqps[0], orig_dip);
820         apic_set_group(irqps[0], orig_cpu);
821         apic_set_group(irqps[0], num_vectors);
822         apic_set_group(irqps[0], vector);
823     }
824
825     /* Set up interrupt table */
826     if (num_vectors > 1) {
827         apic_set_group(irqps[0], new_cpu);
828         apic_set_group(irqps[0], orig_dip);
829         apic_set_group(irqps[0], orig_cpu);
830         apic_set_group(irqps[0], num_vectors);
831         apic_set_group(irqps[0], vector);
832     }
833
834     /* Set up interrupt table */
835     if (num_vectors > 1) {
836         apic_set_group(irqps[0], new_cpu);
837         apic_set_group(irqps[0], orig_dip);
838         apic_set_group(irqps[0], orig_cpu);
839         apic_set_group(irqps[0], num_vectors);
840         apic_set_group(irqps[0], vector);
841     }
842
843     /* Set up interrupt table */
844     if (num_vectors > 1) {
845         apic_set_group(irqps[0], new_cpu);
846         apic_set_group(irqps[0], orig_dip);
847         apic_set_group(irqps[0], orig_cpu);
848         apic_set_group(irqps[0], num_vectors);
849         apic_set_group(irqps[0], vector);
850     }
851
852     /* Set up interrupt table */
853     if (num_vectors > 1) {
854         apic_set_group(irqps[0], new_cpu);
855         apic_set_group(irqps[0], orig_dip);
856         apic_set_group(irqps[0], orig_cpu);
857         apic_set_group(irqps[0], num_vectors);
858         apic_set_group(irqps[0], vector);
859     }
860
861     /* Set up interrupt table */
862     if (num_vectors > 1) {
863         apic_set_group(irqps[0], new_cpu);
864         apic_set_group(irqps[0], orig_dip);
865         apic_set_group(irqps[0], orig_cpu);
866         apic_set_group(irqps[0], num_vectors);
867         apic_set_group(irqps[0], vector);
868     }
869
870     /* Set up interrupt table */
871     if (num_vectors > 1) {
872         apic_set_group(irqps[0], new_cpu);
873         apic_set_group(irqps[0], orig_dip);
874         apic_set_group(irqps[0], orig_cpu);
875         apic_set_group(irqps[0], num_vectors);
876         apic_set_group(irqps[0], vector);
877     }
878
879     /* Set up interrupt table */
880     if (num_vectors > 1) {
881         apic_set_group(irqps[0], new_cpu);
882         apic_set_group(irqps[0], orig_dip);
883         apic_set_group(irqps[0], orig_cpu);
884         apic_set_group(irqps[0], num_vectors);
885         apic_set_group(irqps[0], vector);
886     }
887
888     /* Set up interrupt table */
889     if (num_vectors > 1) {
890         apic_set_group(irqps[0], new_cpu);
891         apic_set_group(irqps[0], orig_dip);
892         apic_set_group(irqps[0], orig_cpu);
893         apic_set_group(irqps[0], num_vectors);
894         apic_set_group(irqps[0], vector);
895     }
896
897     /* Set up interrupt table */
898     if (num_vectors > 1) {
899         apic_set_group(irqps[0], new_cpu);
900         apic_set_group(irqps[0], orig_dip);
901         apic_set_group(irqps[0], orig_cpu);
902         apic_set_group(irqps[0], num_vectors);
903         apic_set_group(irqps[0], vector);
904     }
905
906     /* Set up interrupt table */
907     if (num_vectors > 1) {
908         apic_set_group(irqps[0], new_cpu);
909         apic_set_group(irqps[0], orig_dip);
910         apic_set_group(irqps[0], orig_cpu);
911         apic_set_group(irqps[0], num_vectors);
912         apic_set_group(irqps[0], vector);
913     }
914
915     /* Set up interrupt table */
916     if (num_vectors > 1) {
917         apic_set_group(irqps[0], new_cpu);
918         apic_set_group(irqps[0], orig_dip);
919         apic_set_group(irqps[0], orig_cpu);
920         apic_set_group(irqps[0], num_vectors);
921         apic_set_group(irqps[0], vector);
922     }
923
924     /* Set up interrupt table */
925     if (num_vectors > 1) {
926         apic_set_group(irqps[0], new_cpu);
927         apic_set_group(irqps[0], orig_dip);
928         apic_set_group(irqps[0], orig_cpu);
929         apic_set_group(irqps[0], num_vectors);
930         apic_set_group(irqps[0], vector);
931     }
932
933     /* Set up interrupt table */
934     if (num_vectors > 1) {
935         apic_set_group(irqps[0], new_cpu);
936         apic_set_group(irqps[0], orig_dip);
937         apic_set_group(irqps[0], orig_cpu);
938         apic_set_group(irqps[0], num_vectors);
939         apic_set_group(irqps[0], vector);
940     }
941
942     /* Set up interrupt table */
943     if (num_vectors > 1) {
944         apic_set_group(irqps[0], new_cpu);
945         apic_set_group(irqps[0], orig_dip);
946         apic_set_group(irqps[0], orig_cpu);
947         apic_set_group(irqps[0], num_vectors);
948         apic_set_group(irqps[0], vector);
949     }
950
951     /* Set up interrupt table */
952     if (num_vectors > 1) {
953         apic_set_group(irqps[0], new_cpu);
954         apic_set_group(irqps[0], orig_dip);
955         apic_set_group(irqps[0], orig_cpu);
956         apic_set_group(irqps[0], num_vectors);
957         apic_set_group(irqps[0], vector);
958     }
959
960     /* Set up interrupt table */
961     if (num_vectors > 1) {
962         apic_set_group(irqps[0], new_cpu);
963         apic_set_group(irqps[0], orig_dip);
964         apic_set_group(irqps[0], orig_cpu);
965         apic_set_group(irqps[0], num_vectors);
966         apic_set_group(irqps[0], vector);
967     }
968
969     /* Set up interrupt table */
970     if (num_vectors > 1) {
971         apic_set_group(irqps[0], new_cpu);
972         apic_set_group(irqps[0], orig_dip);
973         apic_set_group(irqps[0], orig_cpu);
974         apic_set_group(irqps[0], num_vectors);
975         apic_set_group(irqps[0], vector);
976     }
977
978     /* Set up interrupt table */
979     if (num_vectors > 1) {
980         apic_set_group(irqps[0], new_cpu);
981         apic_set_group(irqps[0], orig_dip);
982         apic_set_group(irqps[0], orig_cpu);
983        
```

```

551         goto set_grp_intr_done;
552     }
553     DDI_INTR_IMPLDBG((CE_CONT, "set_grp: num intrs in grp: %d\n",
554                       num_vectors));
555
556     ASSERT((num_vectors + vector) < APIC_MAX_VECTOR);
557
558     *result = EIO;
559
560     /*
561      * All IRQ entries in the table for the given device will be not
562      * shared. Since they are not shared, the dip in the table will
563      * be true to the device of interest.
564      */
565     for (i = 1; i < num_vectors; i++) {
566         irqps[i] = apic_irq_table[apic_vector_to_irq[vector + i]];
567         if (irqps[i] == NULL) {
568             mutex_exit(&airq_mutex);
569             goto set_grp_intr_done;
570         }
571 #ifdef DEBUG
572         /* Sanity check: CPU and dip is the same for all entries. */
573         if ((irqps[i]->airq_dip != orig_dip) ||
574             (irqps[i]->airq_temp_cpu != orig_cpu)) {
575             mutex_exit(&airq_mutex);
576             DDI_INTR_IMPLDBG((CE_CONT,
577                               "set_grp: cpu or dip for vec 0x%x difft than for "
578                               "vec 0x%x\n", vector, vector + i));
579             DDI_INTR_IMPLDBG((CE_CONT,
580                               "cpu: %d vs %d, dip: 0x%p vs 0x%p\n", orig_cpu,
581                               irqps[i]->airq_temp_cpu, (void *)orig_dip,
582                               (void *)irqps[i]->airq_dip));
583             goto set_grp_intr_done;
584         }
585 #endif /* DEBUG */
586     }
587     mutex_exit(&airq_mutex);
588
589     cap_ptr = i_ddi_get_msi_msix_cap_ptr(orig_dip);
590     handle = i_ddi_get_pci_config_handle(orig_dip);
591     msi_ctrl = pci_config_get16(handle, cap_ptr + PCI_MSI_CTRL);
592
593     /* MSI Per vector masking is supported. */
594     if (msi_ctrl & PCI_MSI_PVM_MASK) {
595         if (msi_ctrl & PCI_MSI_64BIT_MASK)
596             msi_mask_off = cap_ptr + PCI_MSI_64BIT_MASKBITS;
597         else
598             msi_mask_off = cap_ptr + PCI_MSI_32BIT_MASK;
599         msi_pvm = pci_config_get32(handle, msi_mask_off);
600         pci_config_put32(handle, msi_mask_off, (uint32_t)-1);
601         DDI_INTR_IMPLDBG((CE_CONT,
602                           "set_grp: pvm supported. Mask set to 0x%x\n",
603                           pci_config_get32(handle, msi_mask_off)));
604     }
605
606     iflag = intr_clear();
607     lock_set(&apic_ioapic_lock);
608
609     /*
610      * Do the first rebind and check for errors. Apic_rebind_all returns
611      * an error if the CPU is not accepting interrupts. If the first one
612      * succeeds they all will.
613      */
614     if (apic_rebind_all(irqps[0], new_cpu))
615         (void) apic_rebind_all(irqps[0], orig_cpu);
616     else {

```

```

617         irqps[0]->airq_cpu = new_cpu;
618
619         for (i = 1; i < num_vectors; i++) {
620             (void) apic_rebind_all(irqps[i], new_cpu);
621             irqps[i]->airq_cpu = new_cpu;
622         }
623         *result = 0; /* SUCCESS */
624     }
625
626     lock_clear(&apic_ioapic_lock);
627     intr_restore(iflag);
628
629     /* Reenable vectors if per vector masking is supported. */
630     if (msi_ctrl & PCI_MSI_PVM_MASK) {
631         pci_config_put32(handle, msi_mask_off, msi_pvm);
632         DDI_INTR_IMPLDBG((CE_CONT,
633                           "set_grp: pvm supported. Mask restored to 0x%x\n",
634                           pci_config_get32(handle, msi_mask_off)));
635     }
636
637     set_grp_intr_done:
638     if (*result != 0)
639         return (PSM_FAILURE);
640
641     return (PSM_SUCCESS);
642 }
643
644 int
645 apic_get_vector_intr_info(int vecirq, apic_get_intr_t *intr_params_p)
646 {
647     struct autovec *av_dev;
648     uchar_t irqno;
649     uint i;
650     int i;
651     apic_irq_t *irq_p;
652
653     /* Sanity check the vector/irq argument. */
654     ASSERT(vecirq >= 0 || (vecirq <= APIC_MAX_VECTOR));
655
656     mutex_enter(&airq_mutex);
657
658     /*
659      * Convert the vecirq arg to an irq using vector_to_irq table
660      * if the arg is a vector. Pass thru if already an irq.
661      */
662     if ((intr_params_p->avgi_req_flags & PSMGI_INTRBY_VEC) ==
663         PSMGI_INTRBY_VEC)
664         irqno = apic_vector_to_irq(vecirq);
665     else
666         irqno = (uchar_t)vecirq;
667     irqno = vecirq;
668
669     irq_p = apic_irq_table[irqno];
670
671     if ((irq_p == NULL) ||
672         ((irq_p->airq_mps_intr_index != RESERVE_INDEX) &&
673          ((irq_p->airq_temp_cpu == IRQ_UNBOUND) ||
674           (irq_p->airq_temp_cpu == IRQ_UNINIT)))) {
675         mutex_exit(&airq_mutex);
676         return (PSM_FAILURE);
677     }
678
679     if (intr_params_p->avgi_req_flags & PSMGI_REQ_CPUID) {
680
681         /* Get the (temp) cpu from apic_irq table, indexed by irq. */
682         intr_params_p->avgi_cpu_id = irq_p->airq_temp_cpu;

```

```

682     /* Return user bound info for intrd. */
683     if (intr_params_p->avgi_cpu_id & IRQ_USER_BOUND) {
684         intr_params_p->avgi_cpu_id &= ~IRQ_USER_BOUND;
685         intr_params_p->avgi_cpu_id |= PSMGI_CPU_USER_BOUND;
686     }
687 }
688
689 if (intr_params_p->avgi_req_flags & PSMGI_REQ_VECTOR)
690     intr_params_p->avgi_vector = irq_p->irq_vector;
691
692 if (intr_params_p->avgi_req_flags &
693     (PSMGI_REQ_NUM_DEVS | PSMGI_REQ_GET_DEVS))
694     /* Get number of devices from apic_irq table shared field. */
695     intr_params_p->avgi_num_devs = irq_p->irq_share;
696
697 if (intr_params_p->avgi_req_flags & PSMGI_REQ_GET_DEVS) {
698
699     intr_params_p->avgi_req_flags |= PSMGI_REQ_NUM_DEVS;
700
701     /* Some devices have NULL dip. Don't count these. */
702     if (intr_params_p->avgi_num_devs > 0) {
703         for (i = 0, av_dev = autovect[irqno].avh_link;
704              av_dev; av_dev = av_dev->av_link)
705             if (av_dev->av_vector && av_dev->av_dip)
706                 i++;
707         intr_params_p->avgi_num_devs =
708             (uchar_t)MIN(intr_params_p->avgi_num_devs, i);
709         MIN(intr_params_p->avgi_num_devs, i);
710     }
711
712     /* There are no viable dips to return. */
713     if (intr_params_p->avgi_num_devs == 0)
714         intr_params_p->avgi_dip_list = NULL;
715
716     else { /* Return list of dips */
717
718         /* Allocate space in array for that number of devs. */
719         intr_params_p->avgi_dip_list = kmem_zalloc(
720             intr_params_p->avgi_num_devs *
721             sizeof (dev_info_t *),
722             KM_SLEEP);
723
724         /*
725          * Loop through the device list of the autovect table
726          * filling in the dip array.
727          *
728          * Note that the autovect table may have some special
729          * entries which contain NULL dips. These will be
730          * ignored.
731          */
732         for (i = 0, av_dev = autovect[irqno].avh_link;
733              av_dev; av_dev = av_dev->av_link)
734             if (av_dev->av_vector && av_dev->av_dip)
735                 intr_params_p->avgi_dip_list[i++] =
736                     av_dev->av_dip;
737     }
738
739     mutex_exit(&irq_mutex);
740
741     return (PSM_SUCCESS);
742 }

```

unchanged portion omitted

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

```
*****
2023 Thu Sep 7 15:25:33 2017
new/usr/src/uts/i86pc/pcplusmp/Makefile
8626 make pcplusmp and apix warning-free
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
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 2017, Joyent, Inc.
27 # Copyright 2016, 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)
```

1

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

```
59 LINT_TARGET      = $(MODULE).lint
60 INSTALL_TARGET   = $(BINARY) $(ROOTMODULE)
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 CERRWARN        += -_gcc=-Wno-unused-function
70 #
73 # For now, disable these lint checks; maintainers should endeavor
74 # to investigate and remove these for maximum lint coverage.
75 # Please do not carry these forward to new Makefiles.
76 #
77 LINTTAGS         += -erroff=E_BAD_PTR_CAST_ALIGN
78 LINTTAGS         += -erroff=E_SUPPRESSION_DIRECTIVE_UNUSED
79 LINTTAGS         += -erroff=E_STATIC_UNUSED
80 LINTTAGS         += -erroff=E_ASSIGN_NARROW_CONV
82 CERRWARN        += -_gcc=-Wno-uninitialized
84 #
71 # Default build targets.
72 #
73 .KEEP_STATE:
75 def:           $(DEF_DEPS)
77 all:           $(ALL_DEPS)
79 clean:         $(CLEAN_DEPS)
81 clobber:       $(CLOBBER_DEPS)
83 lint:          $(LINT_DEPS)
85 modlintlib:    $(MODLINTLIB_DEPS)
87 clean.lint:    $(CLEAN_LINT_DEPS)
89 install:       $(INSTALL_DEPS)
91 #
92 # Include common targets.
93 #
94 include $(UTSBASE)/i86pc/Makefile.targ
```

2

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

1166 Thu Sep 7 15:25:33 2017

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

8626 make pcplusmp and apix warning-free

Reviewed by: Robert Mustacchi <rm@joyent.com>

Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>

```
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) 2010, Oracle and/or its affiliates. All rights reserved.
23 * Copyright 2017 Joyent, Inc.
24 */
25
26 #ifndef __SYS_APIX_APIX_H
27 #define __SYS_APIX_APIX_H
28
29 #include <sys/note.h>
30 #include <sys/avintr.h>
31 #include <sys/traptrace.h>
32 #include <sys/apic.h>
33 #include <sys/apic_common.h>
34 #include <sys/apic_timer.h>
35
36 #ifdef __cplusplus
37 extern "C" {
38 #endif
39
40 #ifdef DEBUG
41 #ifndef TRAPTRACE
42 #define TRAPTRACE
43 #endif
44 #endif
45
46 #define APIX_NAME "apix"
47
48 #define APIX_NVECTOR 256 /* max number of per-cpu vectors */
49 #define APIX_NIRQ 256 /* maximum number of IRQs */
50 #define APIX_INVALID_VECT 0 /* invalid vector */
51
52 /* vector type */
53 #define APIX_TYPE_FIXED DDI_INTR_TYPE_FIXED /* 1 */
54 #define APIX_TYPE_MSI DDI_INTR_TYPE_MSI /* 2 */
55 #define APIX_TYPE_MSIX DDI_INTR_TYPE_MSIX /* 4 */
56 #define APIX_TYPE_IPI 8
57
58 /* vector states */
59 enum {
```

1

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

```
60     APIX_STATE_FREED = 0,
61     APIX_STATE_OBSOLETED, /* 1 */
62     APIX_STATE_ALLOCED, /* 2 */
63     APIX_STATE_ENABLED, /* 3 */
64     APIX_STATE_DISABLED /* 4 */
65 };
```

unchanged portion omitted

```
271 extern struct apix_rebind_info apix_rebindinfo;
```

```
273 #define APIX_SET_REBIND_INFO(_ovp, _nvp) \
274     if (((_ovp)->v_flags & APIX_VECT_MASKABLE) == 0) { \
275         apix_rebindinfo.i_pri = (_ovp)->v_pri; \
276         apix_rebindinfo.i_old_cpuid = (_ovp)->v_cpuid; \
277         apix_rebindinfo.i_old_av = (_ovp)->v_autovect; \
278         apix_rebindinfo.i_new_cpuid = (_nvp)->v_cpuid; \
279         apix_rebindinfo.i_new_av = (_nvp)->v_autovect; \
280         apix_rebindinfo.i_go = 1; \
281     }
```

```
283 #define APIX_CLR_REBIND_INFO() \
284     apix_rebindinfo.i_go = 0
```

```
286 #define APIX_IS_FAKE_INTR(_vector) \
287     (apix_rebindinfo.i_go && (_vector) == APIX_RESV_VECTOR)
```

```
289 #define APIX_DO_FAKE_INTR(_cpu, _vector) \
290     if (APIX_IS_FAKE_INTR(_vector)) { \
291         struct autovec *tp = NULL; \
292         struct autovec *tp1; \
293         if ((_cpu) == apix_rebindinfo.i_old_cpuid) \
294             tp = apix_rebindinfo.i_old_av; \
295         else if ((_cpu) == apix_rebindinfo.i_new_cpuid) \
296             tp = apix_rebindinfo.i_new_av; \
297         ASSERT(tp != NULL); \
298         if (tp->av_vector != NULL && \
299             (tp->av_flags & AV_PENTRY_PEND) == 0) { \
300             tp->av_flags |= AV_PENTRY_PEND; \
301             apix_insert_pending_av(apixs[(_cpu)], tp, \
302             tp->av_prilevel); \
303             apixs[(_cpu)]->x_intr_pending |= \
304                 (1 << tp->av_prilevel); \
305         } \
306     }
```

```
307 extern int apix_add_avintr(void *intr_id, int ipl, avfunc xxintr, char *name,
308     int vector, caddr_t arg1, caddr_t arg2, uint64_t *ticks, dev_info_t *dip);
309 extern void apix_rem_avintr(void *intr_id, int ipl, avfunc xxintr,
310     int virt_vect);
```

```
312 extern uint32_t apix_bind_cpu_locked(dev_info_t *dip);
313 extern apix_vector_t *apix_rebind(apix_vector_t *vecp, processorid_t tocpu,
314     int count);
```

```
316 extern uchar_t apix_alloc_ipi(int ipl);
317 extern apix_vector_t *apix_alloc_intx(dev_info_t *dip, int inum, int irqno);
318 extern int apix_alloc_msi(dev_info_t *dip, int inum, int count, int behavior);
319 extern int apix_alloc_msix(dev_info_t *dip, int inum, int count, int behavior);
320 extern void apix_free_vectors(dev_info_t *dip, int inum, int count, int type);
321 extern void apix_enable_vector(apix_vector_t *vecp);
322 extern void apix_disable_vector(apix_vector_t *vecp);
323 extern int apix_obsolete_vector(apix_vector_t *vecp);
324 extern int apix_find_cont_vector_oncpu(uint32_t cpuid, int count);
```

```
326 extern void apix_set_dev_map(apix_vector_t *vecp, dev_info_t *dip, int inum);
327 extern apix_vector_t *apix_get_dev_map(dev_info_t *dip, int inum, int type);
```

2

```
328 extern apix_vector_t *apix_setup_io_intr(apix_vector_t *vecp);
329 extern void ioapix_init_intr(int mask_apic);
330 extern int apix_get_min_dev_inum(dev_info_t *dip, int type);
331 extern int apix_get_max_dev_inum(dev_info_t *dip, int type);

333 /*
334 * apix.c
335 */
336 extern int apix_addspl(int virtvec, int ipl, int min_ipl, int max_ipl);
337 extern int apix_delspl(int virtvec, int ipl, int min_ipl, int max_ipl);
338 extern void apix_intx_set_vector(int irqno, uint32_t cpuid, uchar_t vector);
339 extern apix_vector_t *apix_intx_get_vector(int irqno);
340 extern void apix_intx_enable(int irqno);
341 extern void apix_intx_disable(int irqno);
342 extern void apix_intx_free(int irqno);
343 extern int apix_intx_rebind(int irqno, processorid_t cpuid, uchar_t vector);
344 extern apix_vector_t *apix_set_cpu(apix_vector_t *vecp, int new_cpu,
345     int *result);
346 extern apix_vector_t *apix_grp_set_cpu(apix_vector_t *vecp, int new_cpu,
347     int *result);
348 extern void apix_level_intr_pre_eoi(int irq);
349 extern void apix_level_intr_post_dispatch(int irq);

351 #ifdef __cplusplus
352 }
```

unchanged portion omitted