

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

1

```
*****
1985 Fri Apr 25 16:07:54 2014
new/usr/src/uts/i86pc/apix/Makefile
4804 apix & pcplusmp are nearly warning free already
Tentatively Reviewed by: Robert Mustacchi <rm@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 #
26 # This makefile drives the production of the pcplusmp "mach"
27 # kernel module.
28 #
29 # pcplusmp implementation architecture dependent
30 #
31 #
32 #
33 # Path to the base of the uts directory tree (usually /usr/src/uts).
34 #
35 UTSBASE = ../../

37 #
38 # Define the module and object file sets.
39 #
40 MODULE = apix
41 OBJECTS = $(APIX_OBJS:%=$(OBJS_DIR)/%)
42 LINTS = $(APIX_OBJS:%.o=$(LINTS_DIR)/%.ln)
43 ROOTMODULE = $(ROOT_PSM_MACH_DIR)/$(MODULE)

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

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

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

60 #
```

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

2

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

65 CERRWARN += -_gcc=-Wno-uninitialized
66 CERRWARN += -_gcc=-Wno-unused-variable
67 CERRWARN += -_gcc=-Wno-unused-function
68 CERRWARN += -_gcc=-Wno-empty-body

67 #
68 # Default build targets.
69 #
70 .KEEP_STATE:

72 def: $(DEF_DEPS)

74 all: $(ALL_DEPS)

76 clean: $(CLEAN_DEPS)

78 clobber: $(CLOBBER_DEPS)

80 lint: $(LINT_DEPS)

82 modlintlib: $(MODLINTLIB_DEPS)

84 clean.lint: $(CLEAN_LINT_DEPS)

86 install: $(INSTALL_DEPS)

88 #
89 # Include common targets.
90 #
91 include $(UTSBASE)/i86pc/Makefile.targ
```

```

*****
39524 Fri Apr 25 16:07:54 2014
new/usr/src/uts/i86pc/io/hpet_acpi.c
4804 apix & pcplusmp are nearly warning free already
Tentatively Reviewed by: Robert Mustacchi <rm@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) 2009, 2010, Oracle and/or its affiliates. All rights reserved.
23 */

25 #include <sys/hpet_acpi.h>
26 #include <sys/hpet.h>
27 #include <sys/bitmap.h>
28 #include <sys/inttypes.h>
29 #include <sys/time.h>
30 #include <sys/sunddi.h>
31 #include <sys/ksynch.h>
32 #include <sys/apic.h>
33 #include <sys/callb.h>
34 #include <sys/clock.h>
35 #include <sys/archsystem.h>
36 #include <sys/cpart.h>

38 static int hpet_init_proxy(int *hpet_vect, iflag_t *hpet_flags);
39 static boolean_t hpet_install_proxy(void);
40 static boolean_t hpet_callback(int code);
41 static boolean_t hpet_cpr(int code);
42 static boolean_t hpet_resume(void);
43 static void hpet_cst_callback(uint32_t code);
44 static boolean_t hpet_deep_idle_config(int code);
45 static int hpet_validate_table(ACPI_TABLE_HPETS *hpet_table);
46 static boolean_t hpet_checksum_table(unsigned char *table, unsigned int len);
47 static void *hpet_memory_map(ACPI_TABLE_HPETS *hpet_table);
48 static int hpet_start_main_counter(hpet_info_t *hip);
49 static int hpet_stop_main_counter(hpet_info_t *hip);
50 static uint64_t hpet_read_main_counter_value(hpet_info_t *hip);
51 static uint64_t hpet_set_leg_rt_cnf(hpet_info_t *hip, uint32_t new_value);
52 static uint64_t hpet_read_gen_cap(hpet_info_t *hip);
53 static uint64_t hpet_read_gen_config(hpet_info_t *hip);
54 static uint64_t hpet_read_gen_intrpt_stat(hpet_info_t *hip);
55 static uint64_t hpet_read_timer_N_config(hpet_info_t *hip, uint_t n);
56 static hpet_TN_conf_cap_t hpet_convert_timer_N_config(uint64_t conf);
57 /* LINTED E_STATIC_UNUSED */
58 static uint64_t hpet_read_timer_N_comp(hpet_info_t *hip, uint_t n);
59 /* LINTED E_STATIC_UNUSED */
60 static void hpet_write_gen_cap(hpet_info_t *hip, uint64_t l);

```

```

57 static void hpet_write_gen_config(hpet_info_t *hip, uint64_t l);
58 static void hpet_write_gen_intrpt_stat(hpet_info_t *hip, uint64_t l);
59 static void hpet_write_timer_N_config(hpet_info_t *hip, uint_t n, uint64_t l);
60 static void hpet_write_timer_N_comp(hpet_info_t *hip, uint_t n, uint64_t l);
61 static void hpet_disable_timer(hpet_info_t *hip, uint32_t timer_n);
62 static void hpet_enable_timer(hpet_info_t *hip, uint32_t timer_n);
63 /* LINTED E_STATIC_UNUSED */
64 static void hpet_write_main_counter_value(hpet_info_t *hip, uint64_t l);
65 static int hpet_get_IOAPIC_intr_capable_timer(hpet_info_t *hip);
66 static int hpet_timer_available(uint32_t allocated_timers, uint32_t n);
67 static void hpet_timer_alloc(uint32_t *allocated_timers, uint32_t n);
68 static void hpet_timer_set_up(hpet_info_t *hip, uint32_t timer_n,
69                               uint32_t interrupt);
70 static uint_t hpet_isr(char *arg);
71 static uint32_t hpet_install_interrupt_handler(uint_t (*func)(char *),
72                                                int vector);
73 static void hpet_uninstall_interrupt_handler(void);
74 static void hpet_expire_all(void);
75 static boolean_t hpet_guaranteed_schedule(hrtime_t required_wakeup_time);
76 static boolean_t hpet_use_hpet_timer(hrtime_t *expire);
77 static void hpet_use_lapic_timer(hrtime_t *expire);
78 static void hpet_init_proxy_data(void);

78 /*
79  * hpet_state_lock is used to synchronize disabling/enabling deep c-states
80  * and to synchronize suspend/resume.
81  */
82 static kmutex_t hpet_state_lock;
83 static struct hpet_state {
84     boolean_t proxy_installed; /* CBE proxy interrupt setup */
85     boolean_t cpr; /* currently in CPR */
86     boolean_t cpu_deep_idle; /* user enable/disable */
87     boolean_t uni_cstate; /* disable if only one cstate */
88 } hpet_state = { B_FALSE, B_FALSE, B_TRUE, B_TRUE };
      unchanged portion omitted

522 static uint64_t
523 hpet_read_timer_N_comp(hpet_info_t *hip, uint_t n)
524 {
525     if (hip->timer_n_config[n].size_cap == 1)
526         return (*(uint64_t *)
527                HPET_TIMER_N_COMP_ADDRESS(hip->logical_address, n));
528     else
529         return (*(uint32_t *)
530                HPET_TIMER_N_COMP_ADDRESS(hip->logical_address, n));
531 }

532 static uint64_t
533 hpet_read_main_counter_value(hpet_info_t *hip)
534 {
535     uint64_t value;
536     uint32_t *counter;
537     uint32_t high1, high2, low;

538     counter = (uint32_t *)HPET_MAIN_COUNTER_ADDRESS(hip->logical_address);

539     /*
540      * 32-bit main counters
541      */
542     if (hip->gen_cap.count_size_cap == 0) {
543         value = (uint64_t)*counter;
544         hip->main_counter_value = value;
545         return (value);
546     }

547     /*

```



```

*****
63984 Fri Apr 25 16:07:54 2014
new/usr/src/uts/i86pc/io/mp_platform_misc.c
4804 apix & pcplusmp are nearly warning free already
*****
Tentatively Reviewed by: Robert Mustacchi <rm@joyent.com>
*****
_____unchanged_portion_omitted_____

487 /*
488 * Recompute mask bits for the given interrupt vector.
489 * If there is no interrupt servicing routine for this
490 * vector, this function should disable interrupt vector
491 * from happening at all IPLs. If there are still
492 * handlers using the given vector, this function should
493 * disable the given vector from happening below the lowest
494 * IPL of the remaining handlers.
495 */
496 /*ARGSUSED*/
497 int
498 apic_delspl_common(int irqno, int ipl, int min_ipl, int max_ipl)
499 {
500     uchar_t vector;
501     uint32_t bind_cpu;
502     int intin, irqindex;
503     int ioapic_ix;
504     apic_irq_t *irqptr, *preirqptr, *irqheadptr, *irqp;
505     ulong_t iflag;

507     mutex_enter(&airq_mutex);
508     irqindex = IRQINDEX(irqno);
509     irqptr = preirqptr = irqheadptr = apic_irq_table[irqindex];

511     DDI_INTR_IMPLDBG((CE_CONT, "apic_delspl: dip=0x%p type=%d irqno=0x%x "
512 "vector=0x%x\n", (void *)irqptr->airq_dip,
513 irqptr->airq_mps_intr_index, irqno, irqptr->airq_vector));

515     while (irqptr) {
516         if (VIRTIRQ(irqindex, irqptr->airq_share_id) == irqno)
517             break;
518         preirqptr = irqptr;
519         irqptr = irqptr->airq_next;
520     }
521     ASSERT(irqptr);

523     irqptr->airq_share--;

525     mutex_exit(&airq_mutex);

527     /*
528     * If there are more interrupts at a higher IPL, we don't need
529     * to disable anything.
530     */
531     if (ipl < max_ipl)
532         return (PSM_SUCCESS);

534     /* return if it is not hardware interrupt */
535     if (irqptr->airq_mps_intr_index == RESERVE_INDEX)
536         return (PSM_SUCCESS);

538     if (!apic_picinit_called) {
539         /*
540         * Clear irq_struct. If two devices shared an intpt
541         * line & 1 unloaded before picinit, we are hosed. But, then
542         * we hope the machine survive.
543         */
544         irqptr->airq_mps_intr_index = FREE_INDEX;

```

```

545     irqptr->airq_temp_cpu = IRQ_UNINIT;
546     apic_free_vector(irqptr->airq_vector);
547     return (PSM_SUCCESS);
548 }
549 /*
550 * Downgrade vector to new max_ipl if needed. If we cannot allocate,
551 * use old IPL. Not very elegant, but it should work.
552 */
553 if ((irqptr->airq_ipl != max_ipl) && (max_ipl != PSM_INVALID_IPL) &&
554 !ioapic_mask_workaround[irqptr->airq_ioapicindex]) {
555     apic_irq_t *irqp;
556     if ((vector = apic_allocate_vector(max_ipl, irqno, 1)) {
557         if (vector = apic_allocate_vector(max_ipl, irqno, 1)) {
558             apic_mark_vector(irqheadptr->airq_vector, vector);
559             irqp = irqheadptr;
560             while (irqp) {
561                 irqp->airq_vector = vector;
562                 irqp->airq_ipl = (uchar_t)max_ipl;
563                 if (irqp->airq_temp_cpu != IRQ_UNINIT) {
564                     apic_record_rdt_entry(irqp, irqindex);
565
566                     iflag = intr_clear();
567                     lock_set(&apic_ioapic_lock);

568                     (void) apic_setup_io_intr(irqp,
569 irqindex, B_FALSE);

571                     lock_clear(&apic_ioapic_lock);
572                     intr_restore(iflag);
573                 }
574                 irqp = irqp->airq_next;
575             }
576         }
578     } else if (irqptr->airq_ipl != max_ipl &&
579 max_ipl != PSM_INVALID_IPL &&
580 ioapic_mask_workaround[irqptr->airq_ioapicindex]) {
582     /*
583     * We cannot downgrade the IPL of the vector below the vector's
584     * hardware priority. If we did, it would be possible for a
585     * higher-priority hardware vector to interrupt a CPU running at an IPL
586     * lower than the hardware priority of the interrupting vector (but
587     * higher than the soft IPL of this IRQ). When this happens, we would
588     * then try to drop the IPL BELOW what it was (effectively dropping
589     * below base_spl) which would be potentially catastrophic.
590     *
591     * (e.g. Suppose the hardware vector associated with this IRQ is 0x40
592     * (hardware IPL of 4). Further assume that the old IPL of this IRQ
593     * was 4, but the new IPL is 1. If we forced vector 0x40 to result in
594     * an IPL of 1, it would be possible for the processor to be executing
595     * at IPL 3 and for an interrupt to come in on vector 0x40, interrupting
596     * the currently-executing ISR. When apic_intr_enter consults
597     * apic_irqs[], it will return 1, bringing the IPL of the CPU down to 1
598     * so even though the processor was running at IPL 4, an IPL 1
599     * interrupt will have interrupted it, which must not happen)).
600     *
601     * Effectively, this means that the hardware priority corresponding to
602     * the IRQ's IPL (in apic_ipls[]) cannot be lower than the vector's
603     * hardware priority.
604     *
605     * (In the above example, then, after removal of the IPL 4 device's
606     * interrupt handler, the new IPL will continue to be 4 because the
607     * hardware priority that IPL 1 implies is lower than the hardware
608     * priority of the vector used.)
609     */

```

```

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

```

```

676         * Disable the MSI vector
677         * Make sure we only disable on the last
678         * of the multi-MSI support
679         */
680         if (i_ddi_intr_get_current_nenables(irqptr->airq_dip) == 1) {
681             apic_pci_msi_disable_mode(irqptr->airq_dip,
682                 DDI_INTR_TYPE_MSI);
683         }
684     } else if (irqptr->airq_mps_intr_index == MSIX_INDEX) {
685         /*
686         * Disable the MSI-X vector
687         * needs to clear its mask and addr/data for each MSI-X
688         */
689         apic_pci_msi_unconfigure(irqptr->airq_dip, DDI_INTR_TYPE_MSIX,
690             irqptr->airq_origirq);
691         /*
692         * Make sure we only disable on the last MSI-X
693         */
694         if (i_ddi_intr_get_current_nenables(irqptr->airq_dip) == 1) {
695             apic_pci_msi_disable_mode(irqptr->airq_dip,
696                 DDI_INTR_TYPE_MSIX);
697         }
698     } else {
699         /*
700         * The assumption here is that this is safe, even for
701         * systems with IOAPICs that suffer from the hardware
702         * erratum because all devices have been quiesced before
703         * they unregister their interrupt handlers. If that
704         * assumption turns out to be false, this mask operation
705         * can induce the same erratum result we're trying to
706         * avoid.
707         */
708         ioapic_ix = irqptr->airq_ioapicindex;
709         intin = irqptr->airq_intin_no;
710         ioapic_write(ioapic_ix, APIC_RDT_CMD + 2 * intin, AV_MASK);
711     }
712
713     apic_vt_ops->apic_intrmap_free_entry(&irqptr->airq_intrmap_private);
714
715     /*
716     * This irq entry is the only one in the chain.
717     */
718     if (irqheadptr->airq_next == NULL) {
719         ASSERT(irqheadptr == irqptr);
720         bind_cpu = irqptr->airq_temp_cpu;
721         if (((uint32_t)bind_cpu != IRQ_UNBOUND) &&
722             ((uint32_t)bind_cpu != IRQ_UNINIT)) {
723             ASSERT(apic_cpu_in_range(bind_cpu));
724             if (bind_cpu & IRQ_USER_BOUND) {
725                 /* If hardbound, temp_cpu == cpu */
726                 bind_cpu &= ~IRQ_USER_BOUND;
727                 apic_cpus[bind_cpu].aci_bound--;
728             } else
729                 apic_cpus[bind_cpu].aci_temp_bound--;
730         }
731         irqptr->airq_temp_cpu = IRQ_UNINIT;
732         irqptr->airq_mps_intr_index = FREE_INDEX;
733         lock_clear(&apic_ioapic_lock);
734         intr_restore(iflag);
735         apic_free_vector(irqptr->airq_vector);
736         return (PSM_SUCCESS);
737     }
738
739     /*
740     * If we get here, we are sharing the vector and there are more than
741     * one active irq entries in the chain.

```

```
742     */
743     lock_clear(&apic_ioapic_lock);
744     intr_restore(iflag);

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

757     return (PSM_SUCCESS);
758 }
unchanged portion omitted
```

```
*****
34430 Fri Apr 25 16:07:54 2014
new/usr/src/uts/i86pc/io/pcplusmp/apic.c
4804 apix & pcplusmp are nearly warning free already
Tentatively Reviewed by: Robert Mustacchi <rm@joyent.com>
*****
_unchanged_portion_omitted_
```

```
216 static void *apic_hdlp;
```

```
218 /*
219  * apic_let_idle_redistribute can have the following values:
220  * 0 - If clock decremented it from 1 to 0, clock has to call redistribute.
221  * apic_redistribute_lock prevents multiple idle cpus from redistributing
222  */
223 int apic_num_idle_redistributions = 0;
224 static int apic_let_idle_redistribute = 0;
```

```
218 /* to gather intr data and redistribute */
219 static void apic_redistribute_compute(void);
```

```
221 /*
222  * This is the loadable module wrapper
223  */
```

```
225 int
226 _init(void)
227 {
228     if (apic_coarse_hrttime)
229         apic_ops.psm_gethrtime = &apic_gettime;
230     return (psm_mod_init(&apic_hdlp, &apic_psm_info));
231 }
_unchanged_portion_omitted_
```

```
782 /*
783  * type == -1 indicates it is an internal request. Do not change
784  * resv_vector for these requests
785  */
786 static int
787 apic_get_ipivect(int ipl, int type)
788 {
789     uchar_t vector;
790     int irq;
```

```
792     if ((irq = apic_allocate_irq(APIC_VECTOR(ipl))) != -1) {
793         if ((vector = apic_allocate_vector(ipl, irq, 1)) {
801             if (vector = apic_allocate_vector(ipl, irq, 1)) {
794                 apic_irq_table[irq]->airq_mps_intr_index =
795                     RESERVE_INDEX;
796                 apic_irq_table[irq]->airq_vector = vector;
797                 if (type != -1) {
798                     apic_resv_vector[ipl] = vector;
799                 }
800                 return (irq);
801             }
802         }
803         apic_error |= APIC_ERR_GET_IPIVECT_FAIL;
804         return (-1); /* shouldn't happen */
805 }
_unchanged_portion_omitted_
```

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

1

```
*****
2383 Fri Apr 25 16:07:54 2014
new/usr/src/uts/i86pc/pcplusmp/Makefile
4804 apix & pcplusmp are nearly warning free already
Tentatively Reviewed by: Robert Mustacchi <rm@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 #
28 #
29 # This makefile drives the production of the pcplusmp "mach"
30 # kernel module.
31 #
32 # pcplusmp implementation architecture dependent
33 #
34 #
35 #
36 # Path to the base of the uts directory tree (usually /usr/src/uts).
37 #
38 UTSBASE = ../../
39 #
40 #
41 # Define the module and object file sets.
42 #
43 MODULE = pcplusmp
44 OBJECTS = $(PCPLUSMP_OBJS:%=$(OBJS_DIR)/%)
45 LINTS = $(PCPLUSMP_OBJS:%.o=$(LINTS_DIR)/%.ln)
46 ROOTMODULE = $(ROOT_PSM_MACH_DIR)/$(MODULE)
47 #
48 #
49 # Include common rules.
50 #
51 include $(UTSBASE)/i86pc/Makefile.i86pc
52 #
53 #
54 # Define targets
55 #
56 ALL_TARGET = $(BINARY)
57 LINT_TARGET = $(MODULE).lint
58 INSTALL_TARGET = $(BINARY) $(ROOTMODULE)
59 #
60 DEBUG_FLGS =
```

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

2

```
61 $(NOT_RELEASE_BUILD)DEBUG_DEFS += $(DEBUG_FLGS)
62 #
63 #
64 # Depends on ACPI CA interpreter
65 #
66 LDFLAGS += -dy -N misc/acpica
67 #
68 CFLAGS += -Dbug1157974 -Dbug1155030
69 #
70 #
71 # For now, disable these lint checks; maintainers should endeavor
72 # to investigate and remove these for maximum lint coverage.
73 # Please do not carry these forward to new Makefiles.
74 #
75 LINTTAGS += -erroff=E_BAD_PTR_CAST_ALIGN
76 LINTTAGS += -erroff=E_SUPPRESSION_DIRECTIVE_UNUSED
77 #
78 LINTTAGS += -erroff=E_STATIC_UNUSED
79 #
80 LINTTAGS += -erroff=E_ASSIGN_NARROW_CONV
81 #
82 CERRWARN += -_gcc=-Wno-parentheses
83 CERRWARN += -_gcc=-Wno-uninitialized
84 CERRWARN += -_gcc=-Wno-unused-function
85 CERRWARN += -_gcc=-Wno-unused-variable
86 CERRWARN += -_gcc=-Wno-empty-body
87 #
88 #
89 # Default build targets.
90 #
91 .KEEP_STATE:
92 #
93 def: $(DEF_DEPS)
94 #
95 all: $(ALL_DEPS)
96 #
97 clean: $(CLEAN_DEPS)
98 #
99 clobber: $(CLOBBER_DEPS)
100 #
101 lint: $(LINT_DEPS)
102 #
103 modlintlib: $(MODLINTLIB_DEPS)
104 #
105 clean.lint: $(CLEAN_LINT_DEPS)
106 #
107 install: $(INSTALL_DEPS)
108 #
109 #
110 # Include common targets.
111 #
112 include $(UTSBASE)/i86pc/Makefile.targ
```