

```

*****
48536 Sat Feb 23 00:30:14 2013
new/usr/src/uts/i86pc/io/apix/apix_utils.c
3426 assertion failed: irq < 16 on VMware hardware version 9 (apix related)
Reviewed by: Albert Lee <trisk@nexenta.com>
Reviewed by: Dan McDonald <danmcd@nexenta.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 */

22 /*
23 * Copyright (c) 2010, Oracle and/or its affiliates. All rights reserved.
24 */
25 /*
26 * Copyright (c) 2010, Intel Corporation.
27 * All rights reserved.
28 */
29 /*
30 * Copyright 2013 Nexenta Systems, Inc. All rights reserved.
31 */

33 #include <sys/processor.h>
34 #include <sys/time.h>
35 #include <sys/psm.h>
36 #include <sys/smp_impldefs.h>
37 #include <sys/cram.h>
38 #include <sys/acpi/acpi.h>
39 #include <sys/acpica.h>
40 #include <sys/psm_common.h>
41 #include <sys/pit.h>
42 #include <sys/ddi.h>
43 #include <sys/sunddi.h>
44 #include <sys/ddi_impldefs.h>
45 #include <sys/pci.h>
46 #include <sys/promif.h>
47 #include <sys/x86_archext.h>
48 #include <sys/cpc_impl.h>
49 #include <sys/uadmin.h>
50 #include <sys/panic.h>
51 #include <sys/debug.h>
52 #include <sys/archsystem.h>
53 #include <sys/trap.h>
54 #include <sys/machsystem.h>
55 #include <sys/sysmacros.h>
56 #include <sys/cpuvar.h>
57 #include <sys/rm_platter.h>
58 #include <sys/privregs.h>
59 #include <sys/note.h>

```

```

60 #include <sys/pci_intr_lib.h>
61 #include <sys/spl.h>
62 #include <sys/clock.h>
63 #include <sys/dditypes.h>
64 #include <sys/sunddi.h>
65 #include <sys/x_call.h>
66 #include <sys/reboot.h>
67 #include <sys/apix.h>

69 static int apix_get_avail_vector_oncpu(uint32_t, int, int);
70 static apix_vector_t *apix_init_vector(processorid_t, uchar_t);
71 static void apix_cleanup_vector(apix_vector_t *);
72 static void apix_insert_av(apix_vector_t *, void *, avfunc, caddr_t, caddr_t,
73     uint64_t *, int, dev_info_t *);
74 static void apix_remove_av(apix_vector_t *, struct autovec *);
75 static void apix_clear_dev_map(dev_info_t *, int, int);
76 static boolean_t apix_is_cpu_enabled(processorid_t);
77 static void apix_wait_till_seen(processorid_t, int);

79 #define GET_INTR_INUM(ihdlp) \
80     (((ihdlp) != NULL) ? ((ddi_intr_handle_impl_t *) (ihdlp))->ih_inum : 0)

82 apix_rebind_info_t apix_rebindinfo = {0, 0, 0, NULL, 0, NULL};

84 /*
85  * Allocate IPI
86  *
87  * Return vector number or 0 on error
88  */
89 uchar_t
90 apix_alloc_ipi(int ipl)
91 {
92     apix_vector_t *vecp;
93     uchar_t vector;
94     int cpun;
95     int nproc;

97     APIX_ENTER_CPU_LOCK(0);

99     vector = apix_get_avail_vector_oncpu(0, APIX_IPI_MIN, APIX_IPI_MAX);
100     if (vector == 0) {
101         APIX_LEAVE_CPU_LOCK(0);
102         cmn_err(CE_WARN, "apix: no available IPI\n");
103         apic_error |= APIC_ERR_GET_IPIVECT_FAIL;
104         return (0);
105     }

107     nproc = max(apic_nproc, apic_max_nproc);
108     for (cpun = 0; cpun < nproc; cpun++) {
109         vecp = xv_vector(cpun, vector);
110         if (vecp == NULL) {
111             vecp = kmem_zalloc(sizeof (apix_vector_t), KM_NOSLEEP);
112             if (vecp == NULL) {
113                 cmn_err(CE_WARN, "apix: No memory for ipi");
114                 goto fail;
115             }
116             xv_vector(cpun, vector) = vecp;
117         }
118         vecp->v_state = APIX_STATE_ALLOCED;
119         vecp->v_type = APIX_TYPE_IPI;
120         vecp->v_cpuid = vecp->v_bound_cpuid = cpun;
121         vecp->v_vector = vector;
122         vecp->v_pri = ipl;
123     }
124     APIX_LEAVE_CPU_LOCK(0);
125     return (vector);

```

```

127 fail:
128     while (--cpun >= 0)
129         apix_cleanup_vector(xv_vector(cpun, vector));
130     APIX_LEAVE_CPU_LOCK(0);
131     return (0);
132 }

```

unchanged portion omitted

```

1819 /*
1820  * For interrupts which call add_avintr() before apic is initialized.
1821  * ioapix_setup_intr() will
1822  *   - allocate vector
1823  *   - copy over ISR
1824  */
1825 static void
1826 ioapix_setup_intr(int irqno, iflag_t *flagp)
1827 {
1828     extern struct av_head autovect[];
1829     apix_vector_t *vecp;
1830     apic_irq_t *irqp;
1831     uchar_t ioapicindex, ipin;
1832     ulong_t iflag;
1833     struct autovect *avp;

```

```

1832     irqp = apic_irq_table[irqno];
1833     ioapicindex = acpi_find_ioapic(irqno);
1834     ASSERT(ioapicindex != 0xFF);
1835     ipin = irqno - apic_io_vectbase[ioapicindex];

```

```

1839     mutex_enter(&airq_mutex);
1840     irqp = apic_irq_table[irqno];

```

```

1842     /*
1843     * The irq table entry should not exist unless the interrupts are shared
1844     * In that case, make sure it matches what we would initialize it to.
1845     */
1846     if (irqp != NULL) {
1847         ASSERT(irqp->airq_mps_intr_index == ACPI_INDEX);
1848         if ((irqp != NULL) && (irqp->airq_mps_intr_index == ACPI_INDEX)) {
1849             ASSERT(irqp->airq_intin_no == ipin &&
1850                 irqp->airq_ioapicindex == ioapicindex);
1851             vecp = xv_vector(irqp->airq_cpu, irqp->airq_vector);
1852             ASSERT(!IS_VECT_FREE(vecp));
1853             mutex_exit(&airq_mutex);
1854         } else {
1855             irqp = kmem_zalloc(sizeof(apic_irq_t), KM_SLEEP);
1856             vecp = apix_alloc_intx(NULL, 0, irqno);

```

```

1856             irqp->airq_cpu = IRQ_UNINIT;
1857             irqp->airq_origirq = (uchar_t)irqno;
1858             irqp = apic_irq_table[irqno];
1859             irqp->airq_mps_intr_index = ACPI_INDEX;
1860             irqp->airq_ioapicindex = ioapicindex;
1861             irqp->airq_intin_no = ipin;
1862             irqp->airq_iflag = *flagp;
1863             irqp->airq_share++;

```

```

1864             apic_irq_table[irqno] = irqp;
1865             mutex_exit(&airq_mutex);

```

```

1867             vecp = apix_alloc_intx(NULL, 0, irqno);
1868             apic_record_rdt_entry(irqp, irqno);
1869         }

```

```

1870     /* copy over autovect */

```

```

1871     for (avp = autovect[irqno].avh_link; avp; avp = avp->av_link)
1872         apix_insert_av(vecp, avp->av_intr_id, avp->av_vector,
1873             avp->av_intarg1, avp->av_intarg2, avp->av_ticksps,
1874             avp->av_prilevel, avp->av_dip);

```

```

1876     /* Program I/O APIC */
1877     iflag = intr_clear();
1878     lock_set(&apix_lock);

```

```

1880     (void) apix_setup_io_intr(vecp);

```

```

1882     lock_clear(&apix_lock);
1883     intr_restore(iflag);

```

```

1885     APIC_VERBOSE_IOAPIC((CE_CONT, "apix: setup ioapic, irqno %x "
1886         "(ioapic %x, ipin %x) is bound to cpu %x, vector %x\n",
1887         irqno, ioapicindex, ipin, irqp->airq_cpu, irqp->airq_vector));
1888 }

```

unchanged portion omitted