

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

1

```
*****
65575 Sun Mar  2 12:33:06 2014
new/usr/src/uts/i86pc/io/apix/apix.c
4664 CPU->cpu_pri_data hasn't been used for years
*****
_____unchanged_portion_omitted_____

350 #if defined(__amd64)
351 static unsigned char dummy_cpu_pri[MAXIPL + 1] = {
352     0, 0, 0, 0, 0, 0, 0, 0, 0,
353     0, 0, 0, 0, 0, 0, 0, 0, 0
354 };
355 #endif

350 static void
351 apix_init()
352 {
353     extern void (*do_interrupt_common)(struct regs *, trap_trace_rec_t *);
354
355     APIC_VERBOSE(INIT, (CE_CONT, "apix: psm_softinit\n"));
356
357     do_interrupt_common = apix_do_interrupt;
358     addintr = apix_add_avintr;
359     remintr = apix_rem_avintr;
360     get_pending_spl = apix_get_pending_spl;
361     get_intr_handler = apix_get_intr_handler;
362     psm_get_localapicid = apic_get_localapicid;
363     psm_get_ioapicid = apic_get_ioapicid;
364
365     apix_softinit();
366
367 #if !defined(__amd64)
368 #if defined(__amd64)
369     CPU->cpu_pri_data = dummy_cpu_pri;
370 #else
371     if (cpuid_have_cr8access(CPU))
372         apic_have_32bit_cr8 = 1;
373 #endif
374 #endif
375 #endif /* __amd64 */

372 /*
373  * Initialize IRM pool parameters
374  */
375 if (irm_enable) {
376     int i;
377     int lowest_irq;
378     int highest_irq;
379
380     /* number of CPUs present */
381     apix_irminfo.apix_ncpus = apic_nproc;
382     /* total number of entries in all of the IOAPICs present */
383     lowest_irq = apic_io_vectbase[0];
384     highest_irq = apic_io_vectend[0];
385     for (i = 1; i < apic_io_max; i++) {
386         if (apic_io_vectbase[i] < lowest_irq)
387             lowest_irq = apic_io_vectbase[i];
388         if (apic_io_vectend[i] > highest_irq)
389             highest_irq = apic_io_vectend[i];
390     }
391     apix_irminfo.apix_ioapic_max_vectors =
392         highest_irq - lowest_irq + 1;
393     /*
394      * Number of available per-CPU vectors excluding
395      * reserved vectors for Dtrace, int80, system-call,
396      * fast-trap, etc.
397      */

```

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

2

```
398     apix_irminfo.apix_per_cpu_vectors = APIX_NAVINTR -
399         APIX_SW_RESERVED_VECTORS;
400
401     /* Number of vectors (pre) allocated (SCI and HPET) */
402     apix_irminfo.apix_vectors_allocated = 0;
403     if (apic_hpet_vect != -1)
404         apix_irminfo.apix_vectors_allocated++;
405     if (apic_sci_vect != -1)
406         apix_irminfo.apix_vectors_allocated++;
407     }
408 }
_____unchanged_portion_omitted_____

```

```

*****
35863 Sun Mar  2 12:33:06 2014
new/usr/src/uts/i86pc/io/pcplusmp/apic.c
4664 CPU->cpu_pri_data hasn't been used for years
*****
_____unchanged_portion_omitted_____
137 /*
138  * The ipl of an ISR at vector X is apic_vectortoipl[X>>4]
139  * NOTE that this is vector as passed into intr_enter which is
140  * programmed vector - 0x20 (APIC_BASE_VECT)
141  */

143 uchar_t apic_ipltopri[MAXIPL + 1]; /* unix ipl to apic pri */
144 /* The taskpri to be programmed into apic to mask given ipl */

146 #if defined(__amd64)
147 static unsigned char dummy_cpu_pri[MAXIPL + 1];
148 #endif

146 /*
147  * Correlation of the hardware vector to the IPL in use, initialized
148  * from apic_vectortoipl[] in apic_init(). The final IPLs may not correlate
149  * to the IPLs in apic_vectortoipl on some systems that share interrupt lines
150  * connected to errata-stricken IOAPICs
151  */
152 uchar_t apic_ipls[APIC_AVAIL_VECTOR];

154 /*
155  * Patchable global variables.
156  */
157 int apic_enable_hwssoftint = 0; /* 0 - disable, 1 - enable */
158 int apic_enable_bind_log = 1; /* 1 - display interrupt binding log */

160 /*
161  * Local static data
162  */
163 static struct psm_ops apic_ops = {
164     apic_probe,

166     apic_init,
167     apic_picinit,
168     apic_intr_enter,
169     apic_intr_exit,
170     apic_setspl,
171     apic_addspl,
172     apic_delspl,
173     apic_disable_intr,
174     apic_enable_intr,
175     (int (*)(int))NULL, /* psm_softlvl_to_irq */
176     (void (*)(int))NULL, /* psm_set_softintr */

178     apic_set_idlecpcu,
179     apic_unset_idlecpcu,

181     apic_clkinit,
182     apic_getclkirq,
183     (void (*)(void))NULL, /* psm_hrttimeinit */
184     apic_gethrtime,

186     apic_get_next_processorid,
187     apic_cpu_start,
188     apic_post_cpu_start,
189     apic_shutdown,
190     apic_get_ipivect,
191     apic_send_ipi,

```

```

193     (int (*)(dev_info_t *, int))NULL, /* psm_translate_irq */
194     (void (*)(int, char *))NULL, /* psm_notify_error */
195     (void (*)(int))NULL, /* psm_notify_func */
196     apic_timer_reprogram,
197     apic_timer_enable,
198     apic_timer_disable,
199     apic_post_cyclic_setup,
200     apic_preshutdown,
201     apic_intr_ops, /* Advanced DDI Interrupt framework */
202     apic_state, /* save, restore apic state for S3 */
203     apic_cpu_ops, /* CPU control interface. */
204 };
_____unchanged_portion_omitted_____

274 void
275 apic_init(void)
276 {
277     int i;
278     int j = 1;

280     psm_get_ioapicid = apic_get_ioapicid;
281     psm_get_localapicid = apic_get_localapicid;
282     psm_xlate_vector_by_irq = apic_xlate_vector_by_irq;

284     apic_ipltopri[0] = APIC_VECTOR_PER_IPL; /* leave 0 for idle */
285     for (i = 0; i < (APIC_AVAIL_VECTOR / APIC_VECTOR_PER_IPL); i++) {
286         if ((i < ((APIC_AVAIL_VECTOR / APIC_VECTOR_PER_IPL) - 1)) &&
287             (apic_vectortoipl[i + 1] == apic_vectortoipl[i]))
288             /* get to highest vector at the same ipl */
289             continue;
290         for (; j <= apic_vectortoipl[i]; j++) {
291             apic_ipltopri[j] = (i << APIC_IPL_SHIFT) +
292                 APIC_BASE_VECT;
293         }
294     }
295     for (; j < MAXIPL + 1; j++)
296         /* fill up any empty ipltopri slots */
297         apic_ipltopri[j] = (i << APIC_IPL_SHIFT) + APIC_BASE_VECT;
298     apic_init_common();

300 #if !defined(__amd64)
301 #if defined(__amd64)
302     CPU->cpu_pri_data = dummy_cpu_pri;
303 #else
304     if (cpuid_have_cr8access(CPU))
305         apic_have_32bit_cr8 = 1;
306 #endif
307 #endif /* __amd64 */
308 }
_____unchanged_portion_omitted_____

```

```

*****
7619 Sun Mar  2 12:33:07 2014
new/usr/src/uts/i86pc/ml/offsets.in
4664 CPU->cpu_pri_data hasn't been used for years
*****
1 \
2 \ Copyright (c) 2004, 2010, Oracle and/or its affiliates. All rights reserved.
3 \ Copyright 2012 Garrett D'Amore <garrett@damore.org>. All rights reserved.
4 \
5 \ CDDL HEADER START
6 \
7 \ The contents of this file are subject to the terms of the
8 \ Common Development and Distribution License (the "License").
9 \ You may not use this file except in compliance with the License.
10 \
11 \ You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
12 \ or http://www.opensolaris.org/os/licensing.
13 \ See the License for the specific language governing permissions
14 \ and limitations under the License.
15 \
16 \ When distributing Covered Code, include this CDDL HEADER in each
17 \ file and include the License file at usr/src/OPENSOLARIS.LICENSE.
18 \ If applicable, add the following below this CDDL HEADER, with the
19 \ fields enclosed by brackets "[]" replaced with your own identifying
20 \ information: Portions Copyright [yyyy] [name of copyright owner]
21 \
22 \ CDDL HEADER END
23 \

26 \
27 \ offsets.in: input file to produce assym.h using the ctfstabs program
28 \

30 #ifndef _GENASSYM
31 #define _GENASSYM
32 #endif

34 #define SIZES  1

36 #include <sys/types.h>
37 #include <sys/bootsvcs.h>
38 #include <sys/systm.h>
39 #include <sys/sysinfo.h>
40 #include <sys/user.h>
41 #include <sys/thread.h>
42 #include <sys/proc.h>
43 #include <sys/cpuvar.h>
44 #include <sys/tss.h>
45 #include <sys/privregs.h>
46 #include <sys/segments.h>
47 #include <sys/devops.h>
48 #include <sys/ddi_impldefs.h>
49 #include <vm/as.h>
50 #include <sys/avintr.h>
51 #include <sys/pic.h>
52 #include <sys/rm_platter.h>
53 #include <sys/stream.h>
54 #include <sys/strsubr.h>
55 #include <sys/sunddi.h>
56 #include <sys/modctl.h>
57 #include <sys/traptrace.h>
58 #include <sys/ontrap.h>
59 #include <sys/lgrp.h>
60 #include <sys/dtrace.h>
61 #include <sys/brand.h>

```

```

62 #include <sys/fastboot.h>
63 #include <sys/cpr_wakecode.h>

65 proc          PROC_SIZE
66     p_link
67     p_next
68     p_child
69     p_sibling
70     p_sig
71     p_flag
72     p_tlist
73     p_as
74     p_lockp
75     p_user
76     p_ldt
77     p_ldt_desc
78     p_model
79     p_pctx
80     p_agenttp
81     p_zone
82     p_brand
83     p_brand_data

85 _kthread      THREAD_SIZE
86     t_pcb
87     t_lock
88     t_lockstat
89     t_lockp
90     t_lock_flush
91     t_kpri_req
92     t_oldspl
93     t_pri
94     t_pil
95     t_lwp
96     t_procp
97     t_link
98     t_state
99     t_mstate
100    t_preempt_lk
101    t_stk
102    t_swap
103    t_lwpchan.lc_wchan
104    t_flag
105    t_ctx
106    t_lofault
107    t_onfault
108    t_ontrap
109    t_cpu
110    t_lpl
111    t_bound_cpu
112    t_intr
113    t_forw
114    t_back
115    t_sig
116    t_tid
117    t_pre_sys
118    t_preempt
119    t_proc_flag
120    t_startpc
121    t_sysnum
122    t_intr_start
123    _tu._ts._t_astflag
124    _tu._ts._t_post_sys
125    _tu._t_post_sys_ast
126    t_copyops

T_LABEL
T_STACK
T_WCHAN
T_FLAGS
T_ASTFLAG
T_POST_SYS
T_POST_SYS_AST

```

```

128 ctxop
129     save_op           CTXOP_SAVE

131 as
132     a_hat

134 user    USIZEBYTES
135     u_comm
136     u_signal

138 _label_t
139     val      LABEL_VAL

141 \#define LABEL_PC      LABEL_VAL
142 \#define LABEL_SP      _CONST(LABEL_VAL + LABEL_VAL_INCR)
143 \#define T_PC          _CONST(T_LABEL + LABEL_PC)
144 \#define T_SP          _CONST(T_LABEL + LABEL_SP)

146 _klwp
147     lwp_thread
148     lwp_procp
149     lwp_brand
150     lwp_eosys
151     lwp_regs
152     lwp_arg
153     lwp_ap
154     lwp_cursig
155     lwp_state
156     lwp_mstate.ms_acct    LWP_MS_ACCT
157     lwp_mstate.ms_prev    LWP_MS_PREV
158     lwp_mstate.ms_start    LWP_MS_START
159     lwp_mstate.ms_state_start LWP_MS_STATE_START
160     lwp_pcb
161     lwp_ru.sysc          LWP_RU_SYSC

163 \#define LWP_ACCT_USER    _CONST(LWP_MS_ACCT + _MUL(LMS_USER, LWP_MS_ACCT_)
164 \#define LWP_ACCT_SYSTEM _CONST(LWP_MS_ACCT + _MUL(LMS_SYSTEM, LWP_MS_ACC

166 fpu_ctx
167     fpu_regs          FPU_CTX_FPU_REGS
168     fpu_flags         FPU_CTX_FPU_FLAGS
169     fpu_xsave_mask    FPU_CTX_FPU_XSAVE_MASK

171 fxsave_state    FXSAVE_STATE_SIZE
172     fx_fsw        FXSAVE_STATE_FSW
173     fx_mxcsr_mask FXSAVE_STATE_MXCSR_MASK

176 autovec        AUTOVECSIZE
177     av_vector
178     av_intarg1
179     av_intarg2
180     av_ticksp
181     av_link
182     av_prilevel
183     av_dip

185 av_head
186     avh_link
187     avh_hi_pri
188     avh_lo_pri

190 cpu
191     cpu_id
192     cpu_flags
193     cpu_self

```

```

194     cpu_thread
195     cpu_thread_lock
196     cpu_kprunrun
197     cpu_lwp
198     cpu_fpowner
199     cpu_idle_thread
200     cpu_intr_thread
201     cpu_intr_actv
202     cpu_base_spl
203     cpu_intr_stack
204     cpu_stats.sys.cpumigrate    CPU_STATS_SYS_CPUMIGRATE
205     cpu_stats.sys.intr          CPU_STATS_SYS_INTR
206     cpu_stats.sys.intrblk       CPU_STATS_SYS_INTRBLK
207     cpu_stats.sys.syscall       CPU_STATS_SYS_SYSCALL
208     cpu_profile_pc
209     cpu_profile_upc
210     cpu_profile_pil
211     cpu_fttrace.ftd_state       CPU_FTRACE_STATE
212     cpu_mstate
213     cpu_intracct

215 \#define CPU_INTR_ACTV_REF    _CONST(CPU_INTR_ACTV + 2)

217 cpu
218     cpu_m.pil_high_start    CPU_PIL_HIGH_START
219     cpu_m.intrstat          CPU_INTRSTAT
220     cpu_m.mcpu_current_hat    CPU_CURRENT_HAT
221     cpu_m.mcpu_gdt          CPU_GDT
222     cpu_m.mcpu_idt          CPU_IDT
223     cpu_m.mcpu_tss          CPU_TSS
224     cpu_m.mcpu_softinfo     CPU_SOFTINFO
225     cpu_m.mcpu_pri          CPU_PRI
226     cpu_m.mcpu_pri_data     CPU_PRI_DATA
226 \#if defined(__xpv)
227     cpu_m.mcpu_vcpu_info     CPU_VCPU_INFO
228 \#endif

231 machcpu
232     mcpu_pri_data    MCPU_PRI_DATA

230 standard_pic
231     c_curmask
232     c_iplmask

234 ddi_dma_impl
235     dmai_rflags
236     dmai_rdipl

238 dev_info
239     devi_ops          DEVI_DEV_OPS
240     devi_bus_ctl
241     devi_bus_dma_ctl
242     devi_bus_dma_allochdl
243     devi_bus_dma_freehdl
244     devi_bus_dma_bindhdl
245     devi_bus_dma_unbindhdl
246     devi_bus_dma_flush
247     devi_bus_dma_win

249 dev_ops
250     devo_bus_ops      DEVI_BUS_OPS

252 bus_ops
253     bus_ctl           OPS_CTL
254     bus_dma_map       OPS_MAP
255     bus_dma_ctl       OPS_MCTL

```

new/usr/src/uts/i86pc/ml/offsets.in

5

```

256 bus_dma_allochdl OPS_ALLOCHDL
257 bus_dma_freehdl OPS_FREEHDL
258 bus_dma_bindhdl OPS_BINDHDL
259 bus_dma_unbindhdl OPS_UNBINDHDL
260 bus_dma_flush OPS_FLUSH
261 bus_dma_win OPS_WIN

263 sysent SYSENT_SIZE SYSENT_SIZE_SHIFT
264 sy_callc
265 sy_flags
266 sy_narg

268 stdata
269 sd_lock

271 queue
272 q_flag
273 q_next
274 q_stream
275 q_syncq
276 q_qinfo

278 qinit
279 qi_putp

281 syncq
282 sq_flags
283 sq_count
284 sq_lock
285 sq_wait

287 rm_platter
288 rm_idt_lim IDTROFF
289 rm_gdt_lim GDTROFF
290 rm_pdbcr CR3OFF
291 rm_cpu CPUNOFF
292 rm_cr4 CR4OFF
293 rm_cpu_halt_code CPUHALTCODEOFF
294 rm_cpu_halted CPUHALTEDOFF

296 ddi_acc_impl
297 ahi_acc_attr ACC_ATTR
298 ahi_get8 ACC_GETB
299 ahi_get16 ACC_GETW
300 ahi_get32 ACC_GETL
301 ahi_get64 ACC_GETLL
302 ahi_put8 ACC_PUTB
303 ahi_put16 ACC_PUTW
304 ahi_put32 ACC_PUTL
305 ahi_put64 ACC_PUTLL
306 ahi_rep_get8 ACC_REP_GETB
307 ahi_rep_get16 ACC_REP_GETW
308 ahi_rep_get32 ACC_REP_GETL
309 ahi_rep_get64 ACC_REP_GETLL
310 ahi_rep_put8 ACC_REP_PUTB
311 ahi_rep_put16 ACC_REP_PUTW
312 ahi_rep_put32 ACC_REP_PUTL
313 ahi_rep_put64 ACC_REP_PUTLL

315 on_trap_data
316 ot_prot
317 ot_trap
318 ot_trampoline
319 ot_jmpbuf
320 ot_prev
321 ot_handle

```

new/usr/src/uts/i86pc/ml/offsets.in

6

```

322 ot_padl

324 trap_trace_ctl_t __TRAPTR_SIZE TRAPTR_SIZE_SHIFT
325 ttc_next TRAPTR_NEXT
326 ttc_first TRAPTR_FIRST
327 ttc_limit TRAPTR_LIMIT

329 trap_trace_rec_t TRAP_ENT_SIZE
330 ttr_cr2
331 ttr_info.idt_entry.vector TTR_VECTOR
332 ttr_info.idt_entry.ipl TTR_IPL
333 ttr_info.idt_entry.spl TTR_SPL
334 ttr_info.idt_entry.pri TTR_PRI
335 ttr_info.gate_entry.sysnum TTR_SYSNUM
336 ttr_marker
337 ttr_stamp
338 ttr_curthread
339 ttr_sdepth
340 ttr_stack

342 lgrp_ld
343 lpl_lgrp_id

345 dtrace_id_t DTRACE_IDSIZE

347 cpu_core CPU_CORE_SIZE CPU_CORE_SHIFT
348 cpuc_dtrace_flags
349 cpuc_dtrace_illval

351 timespec TIMESPEC_SIZE

353 gate_desc GATE_DESC_SIZE

355 descnbr_t DESCNBR_SIZE
356 dtr_limit
357 dtr_base

359 mod_stub_info MODS_SIZE
360 mods_func_adr MODS_INSTFCN
361 mods_errfcn MODS_RETFCN
362 mods_flag

364 \#define TRAP_TSIZE _MUL(TRAP_ENT_SIZE, TRAPTR_NENT)

366 copyops
367 cp_copyin
368 cp_xcopyin
369 cp_copyout
370 cp_xcopyout
371 cp_copyinstr
372 cp_copyoutstr
373 cp_fuword8
374 cp_fuword16
375 cp_fuword32
376 cp_fuword64
377 cp_suword8
378 cp_suword16
379 cp_suword32
380 cp_suword64
381 cp_physio

383 brand
384 b_machops

386 brand_proc_data_t
387 spd_handler

```

```

389 fastboot_file_t
390     fb_va
391     fb_pte_list_va
392     fb_pte_list_pa
393     fb_dest_pa
394     fb_size
395     fb_next_pa
396     fb_sections
397     fb_sectcnt

399 fastboot_section_t
400     fb_sec_offset
401     fb_sec_paddr
402     fb_sec_size
403     fb_sec_bss_size

405 fastboot_info_t
406     fi_files
407     fi_has_pae
408     fi_pagetable_va
409     fi_pagetable_pa
410     fi_last_table_pa
411     fi_new_mbi_pa
412     fi_valid

414 zone
415     zone_brand_data

417 wc_cpu WC_CPU_SIZE
418     wc_retaddr
419     wc_virtaddr
420     wc_cr0
421     wc_cr3
422     wc_cr4
423     wc_cr8
424     wc_fs
425     wc_fsbases
426     wc_gs
427     wc_gsbases
428     wc_kgsbases
429     wc_r8
430     wc_r9
431     wc_r10
432     wc_r11
433     wc_r12
434     wc_r13
435     wc_r14
436     wc_r15
437     wc_rax
438     wc_rbp
439     wc_rbx
440     wc_rcx
441     wc_rdi
442     wc_rdx
443     wc_rsi
444     wc_rsp
445     wc_gdt_limit    WC_GDT
446     wc_gdt_base
447     wc_idt_limit    WC_IDT
448     wc_idt_base
449     wc_tr
450     wc_ldt
451     wc_eflags
452     wc_ebx
453     wc_edi

```

```

454     wc_esi
455     wc_ebp
456     wc_esp
457     wc_esp
458     wc_ss
459     wc_cs
460     wc_ds
461     wc_es
462     wc_cpu_id
463     wc_saved_stack

465 wc_wakecode
466     wc_cpu

```

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

1

```
*****
12687 Sun Mar  2 12:33:07 2014
new/usr/src/uts/i86pc/os/mlsetup.c
4664 CPU->cpu_pri_data hasn't been used for years
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */
21 /*
22 * Copyright (c) 2012 Gary Mills
23 *
24 * Copyright (c) 1993, 2010, Oracle and/or its affiliates. All rights reserved.
25 * Copyright (c) 2011 by Delphix. All rights reserved.
26 */
27 /*
28 * Copyright (c) 2010, Intel Corporation.
29 * All rights reserved.
30 */
32 #include <sys/types.h>
33 #include <sys/sysmacros.h>
34 #include <sys/disp.h>
35 #include <sys/promif.h>
36 #include <sys/clock.h>
37 #include <sys/cpuvar.h>
38 #include <sys/stack.h>
39 #include <vm/as.h>
40 #include <vm/hat.h>
41 #include <sys/reboot.h>
42 #include <sys/avintr.h>
43 #include <sys/vtrace.h>
44 #include <sys/proc.h>
45 #include <sys/thread.h>
46 #include <sys/cpart.h>
47 #include <sys/pset.h>
48 #include <sys/copyops.h>
49 #include <sys/pg.h>
50 #include <sys/disp.h>
51 #include <sys/debug.h>
52 #include <sys/sunddi.h>
53 #include <sys/x86_archext.h>
54 #include <sys/privregs.h>
55 #include <sys/machsystem.h>
56 #include <sys/onttrap.h>
57 #include <sys/bootconf.h>
58 #include <sys/boot_console.h>
59 #include <sys/kdi_machimpl.h>
60 #include <sys/archsystem.h>
61 #include <sys/promif.h>
```

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

2

```
62 #include <sys/pci_cfgspace.h>
63 #ifdef __xpv
64 #include <sys/hypervisor.h>
65 #else
66 #include <sys/xpv_support.h>
67 #endif
69 /*
70 * some globals for patching the result of cpuid
71 * to solve problems w/ creative cpu vendors
72 */
74 extern uint32_t cpuid_feature_ecx_include;
75 extern uint32_t cpuid_feature_ecx_exclude;
76 extern uint32_t cpuid_feature_edx_include;
77 extern uint32_t cpuid_feature_edx_exclude;
79 /*
80 * Dummy spl priority masks
81 */
82 static unsigned char dummy_cpu_pri[MAXIPL + 1] = {
83     0xf, 0xf, 0xf, 0xf, 0xf, 0xf, 0xf, 0xf,
84     0xf, 0xf, 0xf, 0xf, 0xf, 0xf, 0xf, 0xf
85 };
87 /*
88 * Set console mode
89 */
92 static void
93 set_console_mode(uint8_t val)
94 {
95     struct bop_regs rp = {0};
97     rp.eax.byte.ah = 0x0;
98     rp.eax.byte.al = val;
99     rp.ebx.word.bx = 0x0;
101     BOP_DOINT(bootops, 0x10, &rp);
103 }
105 /*
106 * Setup routine called right before main(). Interposing this function
107 * before main() allows us to call it in a machine-independent fashion.
108 */
109 void
110 mlsetup(struct regs *rp)
111 {
112     u_longlong_t prop_value;
113     extern struct classfuncs sys_classfuncs;
114     extern disp_t cpu0_disp;
115     extern char t0stack[];
116     extern int post_fastreboot;
117     extern uint64_t plat_dr_options;
119     ASSERT_STACK_ALIGNED();
121     /*
122      * initialize cpu_self
123      */
124     cpu[0]->cpu_self = cpu[0];
126 #if defined(__xpv)
127     /*
128      * Point at the hypervisor's virtual cpu structure
129      */
130 
```

```

120     cpu[0]->cpu_m.mcpu_vcpcu_info = &HYPERVISOR_shared_info->vcpcu_info[0];
121 #endif

123     /*
124     * Set up dummy cpu_pri_data values till psm spl code is
125     * installed. This allows splx() to work on amd64.
126     */

136     cpu[0]->cpu_pri_data = dummy_cpu_pri;

138     /*
139     * check if we've got special bits to clear or set
140     * when checking cpu features
141     */

128     if (bootprop_getval("cpuid_feature_ecx_include", &prop_value) != 0)
129         cpuid_feature_ecx_include = 0;
130     else
131         cpuid_feature_ecx_include = (uint32_t)prop_value;

133     if (bootprop_getval("cpuid_feature_ecx_exclude", &prop_value) != 0)
134         cpuid_feature_ecx_exclude = 0;
135     else
136         cpuid_feature_ecx_exclude = (uint32_t)prop_value;

138     if (bootprop_getval("cpuid_feature_edx_include", &prop_value) != 0)
139         cpuid_feature_edx_include = 0;
140     else
141         cpuid_feature_edx_include = (uint32_t)prop_value;

143     if (bootprop_getval("cpuid_feature_edx_exclude", &prop_value) != 0)
144         cpuid_feature_edx_exclude = 0;
145     else
146         cpuid_feature_edx_exclude = (uint32_t)prop_value;

148     /*
149     * Initialize idt0, gdt0, ldt0_default, ktss0 and dftss.
150     */
151     init_desctbls();

153     /*
154     * lgrp_init() and possibly cpuid_pass1() need PCI config
155     * space access
156     */
157 #if defined(__xpv)
158     if (DOMAIN_IS_INITDOMAIN(xen_info))
159         pci_cfgspace_init();
160 #else
161     pci_cfgspace_init();
162     /*
163     * Initialize the platform type from CPU 0 to ensure that
164     * determine_platform() is only ever called once.
165     */
166     determine_platform();
167 #endif

169     /*
170     * The first lightweight pass (pass0) through the cpuid data
171     * was done in locore before mlsetup was called. Do the next
172     * pass in C code.
173     *
174     * The x86_featureset is initialized here based on the capabilities
175     * of the boot CPU. Note that if we choose to support CPUs that have
176     * different feature sets (at which point we would almost certainly
177     * want to set the feature bits to correspond to the feature
178     * minimum) this value may be altered.

```

```

179     /*
180     cpuid_pass1(cpu[0], x86_featureset);

182 #if !defined(__xpv)
183     if ((get_hwenv() & HW_XEN_HVM) != 0)
184         xen_hvm_init();

186     /*
187     * Patch the tsc_read routine with appropriate set of instructions,
188     * depending on the processor family and architecture, to read the
189     * time-stamp counter while ensuring no out-of-order execution.
190     * Patch it while the kernel text is still writable.
191     *
192     * Note: tsc_read is not patched for intel processors whose family
193     * is >6 and for amd whose family >f (in case they don't support rdtscp
194     * instruction, unlikely). By default tsc_read will use cpuid for
195     * serialization in such cases. The following code needs to be
196     * revisited if intel processors of family >= f retains the
197     * instruction serialization nature of mfence instruction.
198     * Note: tsc_read is not patched for x86 processors which do
199     * not support "mfence". By default tsc_read will use cpuid for
200     * serialization in such cases.
201     *
202     * The Xen hypervisor does not correctly report whether rdtscp is
203     * supported or not, so we must assume that it is not.
204     */
205     if ((get_hwenv() & HW_XEN_HVM) == 0 &&
206         is_x86_feature(x86_featureset, X86FSET_TSCP))
207         patch_tsc_read(X86_HAVE_TSCP);
208     else if (cpuid_getvendor(CPU) == X86_VENDOR_AMD &&
209             cpuid_getfamily(CPU) <= 0xf &&
210             is_x86_feature(x86_featureset, X86FSET_SSE2))
211         patch_tsc_read(X86_TSC_MFENCE);
212     else if (cpuid_getvendor(CPU) == X86_VENDOR_Intel &&
213             cpuid_getfamily(CPU) <= 6 &&
214             is_x86_feature(x86_featureset, X86FSET_SSE2))
215         patch_tsc_read(X86_TSC_LFENCE);

217 #endif /* !__xpv */

219 #if defined(__i386) && !defined(__xpv)
220     /*
221     * Some i386 processors do not implement the rdtsc instruction,
222     * or at least they do not implement it correctly. Patch them to
223     * return 0.
224     */
225     if (!is_x86_feature(x86_featureset, X86FSET_TSC))
226         patch_tsc_read(X86_NO_TSC);
227 #endif /* __i386 && !__xpv */

229 #if defined(__amd64) && !defined(__xpv)
230     patch_memops(cpuid_getvendor(CPU));
231 #endif /* __amd64 && !__xpv */

233 #if !defined(__xpv)
234     /* XXPV what, if anything, should be dorked with here under xen? */

236     /*
237     * While we're thinking about the TSC, let's set up %cr4 so that
238     * userland can issue rdtsc, and initialize the TSC_AUX value
239     * (the cpuid) for the rdtscp instruction on appropriately
240     * capable hardware.
241     */
242     if (is_x86_feature(x86_featureset, X86FSET_TSC))
243         setcr4(getcr4() & ~CR4_TSD);

```



```

245     if (is_x86_feature(x86_featureset, X86FSET_TSCP))
246         (void) wrmsr(MSR_AMD_TSCAUX, 0);

248     if (is_x86_feature(x86_featureset, X86FSET_DE))
249         setcr4(getcr4() | CR4_DE);
250 #endif /* __xpv */

252     /*
253      * initialize t0
254      */
255     t0.t_stk = (caddr_t)rp - MINFRAME;
256     t0.t_stkbase = t0stk;
257     t0.t_pri = maxclsypri - 3;
258     t0.t_schedflag = TS_LOAD | TS_DONT_SWAP;
259     t0.t_procp = &p0;
260     t0.t_plockp = &p0lock.pl_lock;
261     t0.t_lwp = &lwp0;
262     t0.t_forw = &t0;
263     t0.t_back = &t0;
264     t0.t_next = &t0;
265     t0.t_prev = &t0;
266     t0.t_cpu = cpu[0];
267     t0.t_disp_queue = &cpu0_disp;
268     t0.t_bind_cpu = PBIND_NONE;
269     t0.t_bind_pset = PS_NONE;
270     t0.t_bindflag = (uchar_t)default_binding_mode;
271     t0.t_cpupart = &cp_default;
272     t0.t_clfuncs = &sys_classfuncs.thread;
273     t0.t_copyops = NULL;
274     THREAD_ONPROC(&t0, CPU);

276     lwp0.lwp_thread = &t0;
277     lwp0.lwp_regs = (void *)rp;
278     lwp0.lwp_procp = &p0;
279     t0.t_tid = p0.p_lwpcnt = p0.p_lwprcnt = p0.p_lwpid = 1;

281     p0.p_exec = NULL;
282     p0.p_stat = SRUN;
283     p0.p_flag = SSYS;
284     p0.p_tlist = &t0;
285     p0.p_stksize = 2*PAGESIZE;
286     p0.p_stkpageszc = 0;
287     p0.p_as = &kas;
288     p0.p_lockp = &p0lock;
289     p0.p_brkpageszc = 0;
290     p0.p_tl_lgrpid = LGRP_NONE;
291     p0.p_tr_lgrpid = LGRP_NONE;
292     sigorset(&p0.p_ignore, &ignoredefault);

294     CPU->cpu_thread = &t0;
295     bzero(&cpu0_disp, sizeof (disp_t));
296     CPU->cpu_disp = &cpu0_disp;
297     CPU->cpu_disp->disp_cpu = CPU;
298     CPU->cpu_dispthread = &t0;
299     CPU->cpu_idle_thread = &t0;
300     CPU->cpu_flags = CPU_READY | CPU_RUNNING | CPU_EXISTS | CPU_ENABLE;
301     CPU->cpu_dispatch_pri = t0.t_pri;

303     CPU->cpu_id = 0;

305     CPU->cpu_pri = 12;          /* initial PIL for the boot CPU */

307     /*
308      * The kernel doesn't use LDTs unless a process explicitly requests one.
309      */
310     p0.p_ldt_desc = null_sdesc;

```

```

312     /*
313      * Initialize thread/cpu microstate accounting
314      */
315     init_mstate(&t0, LMS_SYSTEM);
316     init_cpu_mstate(CPU, CMS_SYSTEM);

318     /*
319      * Initialize lists of available and active CPUs.
320      */
321     cpu_list_init(CPU);

323     pg_cpu_bootstrap(CPU);

325     /*
326      * Now that we have taken over the GDT, IDT and have initialized
327      * active CPU list it's time to inform kmdb if present.
328      */
329     if (boothowto & RB_DEBUG)
330         kdi_idt_sync();

332     /*
333      * Explicitly set console to text mode (0x3) if this is a boot
334      * post Fast Reboot, and the console is set to CONS_SCREEN_TEXT.
335      */
336     if (post_fastreboot && boot_console_type(NULL) == CONS_SCREEN_TEXT)
337         set_console_mode(0x3);

339     /*
340      * If requested (boot -d) drop into kmdb.
341      */
342     /* This must be done after cpu_list_init() on the 64-bit kernel
343      * since taking a trap requires that we re-compute gsbase based
344      * on the cpu list.
345      */
346     if (boothowto & RB_DEBUGENTER)
347         kmdb_enter();

349     cpu_vm_data_init(CPU);

351     rp->r_fp = 0;    /* terminate kernel stack traces! */

353     prom_init("kernel", (void *)NULL);

355     /* User-set option overrides firmware value. */
356     if (bootprop_getval(PLAT_DR_OPTIONS_NAME, &prop_value) == 0) {
357         plat_dr_options = (uint64_t)prop_value;
358     }
359 #if defined(__xpv)
360     /* No support of DR operations on xpv */
361     plat_dr_options = 0;
362 #else
363     /* __xpv */
364     /* Flag PLAT_DR_FEATURE_ENABLED should only be set by DR driver. */
365     plat_dr_options &= ~PLAT_DR_FEATURE_ENABLED;
366 #endif
367     /* Only enable CPU/memory DR on 64 bits kernel. */
368     plat_dr_options &= ~PLAT_DR_FEATURE_MEMORY;
369     plat_dr_options &= ~PLAT_DR_FEATURE_CPU;
370 #endif /* __amd64 */
371 #endif /* __xpv */

372     /*
373      * Get value of "plat_dr_physmax" boot option.
374      * It overrides values calculated from MSCT or SRAT table.
375      */
376     if (bootprop_getval(PLAT_DR_PHYSMAX_NAME, &prop_value) == 0) {

```

```
377     plat_dr_physmax = ((uint64_t)prop_value) >> PAGESHIFT;
378 }
380 /* Get value of boot_ncpus. */
381 if (bootprop_getval(BOOT_NCPUS_NAME, &prop_value) != 0) {
382     boot_ncpus = NCPU;
383 } else {
384     boot_ncpus = (int)prop_value;
385     if (boot_ncpus <= 0 || boot_ncpus > NCPU)
386         boot_ncpus = NCPU;
387 }
389 /*
390  * Set max_ncpus and boot_max_ncpus to boot_ncpus if platform doesn't
391  * support CPU DR operations.
392  */
393 if (plat_dr_support_cpu() == 0) {
394     max_ncpus = boot_max_ncpus = boot_ncpus;
395 } else {
396     if (bootprop_getval(PLAT_MAX_NCPUS_NAME, &prop_value) != 0) {
397         max_ncpus = NCPU;
398     } else {
399         max_ncpus = (int)prop_value;
400         if (max_ncpus <= 0 || max_ncpus > NCPU) {
401             max_ncpus = NCPU;
402         }
403         if (boot_ncpus > max_ncpus) {
404             boot_ncpus = max_ncpus;
405         }
406     }
408     if (bootprop_getval(BOOT_MAX_NCPUS_NAME, &prop_value) != 0) {
409         boot_max_ncpus = boot_ncpus;
410     } else {
411         boot_max_ncpus = (int)prop_value;
412         if (boot_max_ncpus <= 0 || boot_max_ncpus > NCPU) {
413             boot_max_ncpus = boot_ncpus;
414         } else if (boot_max_ncpus > max_ncpus) {
415             boot_max_ncpus = max_ncpus;
416         }
417     }
418 }
420 /*
421  * Initialize the lgrp framework
422  */
423 lgrp_init(LGRP_INIT_STAGE1);
425 if (boothowto & RB_HALT) {
426     prom_printf("unix: kernel halted by -h flag\n");
427     prom_enter_mon();
428 }
430 ASSERT_STACK_ALIGNED();
432 /*
433  * Fill out cpu_ucode_info. Update microcode if necessary.
434  */
435 ucode_check(CPU);
437 if (workaround_errata(CPU) != 0)
438     panic("critical workaround(s) missing for boot cpu");
439 }
```

unchanged portion omitted

```

*****
50258 Sun Mar  2 12:33:07 2014
new/usr/src/uts/i86pc/os/mp_startup.c
4664 CPU->cpu_pri_data hasn't been used for years
*****
_____unchanged_portion_omitted_____

246 /*
247  * Multiprocessor initialization.
248  *
249  * Allocate and initialize the cpu structure, TRAPTRACE buffer, and the
250  * startup and idle threads for the specified CPU.
251  * Parameter boot is true for boot time operations and is false for CPU
252  * DR operations.
253  */
254 static struct cpu *
255 mp_cpu_configure_common(int cpun, boolean_t boot)
256 {
257     struct cpu *cp;
258     kthread_id_t tp;
259     caddr_t sp;
260     proc_t *procp;
261 #if !defined(__xpv)
262     extern int idle_cpu_prefer_mwait;
263     extern void cpu_idle_mwait();
264 #endif
265     extern void idle();
266     extern void cpu_idle();

268 #ifdef TRAPTRACE
269     trap_trace_ctl_t *ttc = &trap_trace_ctl[cpun];
270 #endif

272     ASSERT(MUTEX_HELD(&cpu_lock));
273     ASSERT(cpun < NCPU && cpu[cpun] == NULL);

275     if (cpu_free_list == NULL) {
276         cp = kmem_zalloc(sizeof (*cp), KM_SLEEP);
277     } else {
278         cp = cpu_free_list;
279         cpu_free_list = cp->cpu_next_free;
280     }

282     cp->cpu_m.mcpu_istamp = cpun << 16;

284     /* Create per CPU specific threads in the process p0. */
285     procp = &p0;

287     /*
288      * Initialize the dispatcher first.
289      */
290     disp_cpu_init(cp);

292     cpu_vm_data_init(cp);

294     /*
295      * Allocate and initialize the startup thread for this CPU.
296      * Interrupt and process switch stacks get allocated later
297      * when the CPU starts running.
298      */
299     tp = thread_create(NULL, 0, NULL, NULL, 0, procp,
300         TS_STOPPED, maxclsypri);

302     /*
303      * Set state to TS_ONPROC since this thread will start running
304      * as soon as the CPU comes online.

```

```

305     *
306     * All the other fields of the thread structure are setup by
307     * thread_create().
308     */
309     THREAD_ONPROC(tp, cp);
310     tp->t_preempt = 1;
311     tp->t_bound_cpu = cp;
312     tp->t_affinitycnt = 1;
313     tp->t_cpu = cp;
314     tp->t_disp_queue = cp->cpu_disp;

316     /*
317      * Setup thread to start in mp_startup_common.
318      */
319     sp = tp->t_stk;
320     tp->t_sp = (uintptr_t)(sp - MINFRAME);
321 #if defined(__amd64)
322     tp->t_sp -= STACK_ENTRY_ALIGN; /* fake a call */
323 #endif
324     /*
325      * Setup thread start entry point for boot or hotplug.
326      */
327     if (boot) {
328         tp->t_pc = (uintptr_t)mp_startup_boot;
329     } else {
330         tp->t_pc = (uintptr_t)mp_startup_hotplug;
331     }

333     cp->cpu_id = cpun;
334     cp->cpu_self = cp;
335     cp->cpu_thread = tp;
336     cp->cpu_lwp = NULL;
337     cp->cpu_dispthread = tp;
338     cp->cpu_dispatch_pri = DISP_PRIO(tp);

340     /*
341      * cpu_base_spl must be set explicitly here to prevent any blocking
342      * operations in mp_startup_common from causing the spl of the cpu
343      * to drop to 0 (allowing device interrupts before we're ready) in
344      * resume().
345      * cpu_base_spl MUST remain at LOCK_LEVEL until the cpu is CPU_READY.
346      * As an extra bit of security on DEBUG kernels, this is enforced with
347      * an assertion in mp_startup_common() -- before cpu_base_spl is set
348      * to its proper value.
349      */
350     cp->cpu_base_spl = ipltospl(LOCK_LEVEL);

352     /*
353      * Now, initialize per-CPU idle thread for this CPU.
354      */
355     tp = thread_create(NULL, PAGE_SIZE, idle, NULL, 0, procp, TS_ONPROC, -1);

357     cp->cpu_idle_thread = tp;

359     tp->t_preempt = 1;
360     tp->t_bound_cpu = cp;
361     tp->t_affinitycnt = 1;
362     tp->t_cpu = cp;
363     tp->t_disp_queue = cp->cpu_disp;

365     /*
366      * Bootstrap the CPU's PG data
367      */
368     pg_cpu_bootstrap(cp);

370     /*

```

```

371      * Perform CPC initialization on the new CPU.
372      */
373      kcpc_hw_init(cp);

375      /*
376      * Allocate virtual addresses for cpu_caddr1 and cpu_caddr2
377      * for each CPU.
378      */
379      setup_vaddr_for_ppcopy(cp);

381      /*
382      * Allocate page for new GDT and initialize from current GDT.
383      */
384      #if !defined(__lint)
385          ASSERT((sizeof (*cp->cpu_gdt) * NGDT) <= PAGESIZE);
386      #endif
387      cp->cpu_gdt = kmem_zalloc(PAGESIZE, KM_SLEEP);
388      bcopy(CPU->cpu_gdt, cp->cpu_gdt, (sizeof (*cp->cpu_gdt) * NGDT));

390      #if defined(__i386)
391          /*
392          * setup kernel %gs.
393          */
394          set_usegd(&cp->cpu_gdt[GDT_GS], cp, sizeof (struct cpu) - 1, SDT_MEMRWA,
395                  SEL_KPL, 0, 1);
396      #endif

398      /*
399      * If we have more than one node, each cpu gets a copy of IDT
400      * local to its node. If this is a Pentium box, we use cpu 0's
401      * IDT. cpu 0's IDT has been made read-only to workaroud the
402      * cmpxchgl register bug
403      */
404      if (system_hardware.hd_nodes && x86_type != X86_TYPE_P5) {
405      #if !defined(__lint)
406          ASSERT((sizeof (*CPU->cpu_idt) * NIDT) <= PAGESIZE);
407      #endif
408          cp->cpu_idt = kmem_zalloc(PAGESIZE, KM_SLEEP);
409          bcopy(CPU->cpu_idt, cp->cpu_idt, PAGESIZE);
410      } else {
411          cp->cpu_idt = CPU->cpu_idt;
412      }

414      /*
415      * Get interrupt priority data from cpu 0.
416      */
417      cp->cpu_pri_data = CPU->cpu_pri_data;

419      /*
420      * alloc space for cpuid info
421      */
422      cpuid_alloc_space(cp);
423      #if !defined(__xpv)
424      if (is_x86_feature(x86_featureset, X86FSET_MWAIT) &&
425          idle_cpu_prefer_mwait) {
426          cp->cpu_m.mcpu_mwait = cpuid_mwait_alloc(cp);
427          cp->cpu_m.mcpu_idle_cpu = cpu_idle_mwait;
428      } else
429      #endif
430      cp->cpu_m.mcpu_idle_cpu = cpu_idle;

432      init_cpu_info(cp);

434      /*
435      * alloc space for ucode_info
436      */

```

```

432      ucode_alloc_space(cp);
433      xc_init_cpu(cp);
434      hat_cpu_online(cp);

436      #ifdef TRAPTRACE
437          /*
438          * If this is a TRAPTRACE kernel, allocate TRAPTRACE buffers
439          */
440          ttc->ttc_first = (uintptr_t)kmem_zalloc(trap_trace_bufsize, KM_SLEEP);
441          ttc->ttc_next = ttc->ttc_first;
442          ttc->ttc_limit = ttc->ttc_first + trap_trace_bufsize;
443      #endif

445      /*
446      * Record that we have another CPU.
447      */
448      /*
449      * Initialize the interrupt threads for this CPU
450      */
451      cpu_intr_alloc(cp, NINTR_THREADS);

453      cp->cpu_flags = CPU_OFFLINE | CPU QUIESCED | CPU_POWEROFF;
454      cpu_set_state(cp);

456      /*
457      * Add CPU to list of available CPUs. It'll be on the active list
458      * after mp_startup_common().
459      */
460      cpu_add_unit(cp);

462      return (cp);
463      }
_____unchanged_portion_omitted_____

```

```

*****
5341 Sun Mar  2 12:33:07 2014
new/usr/src/uts/i86pc/sys/machcpuvar.h
4664 CPU->cpu_pri_data hasn't been used for years
*****
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 2009 Sun Microsystems, Inc. All rights reserved.
23 * Use is subject to license terms.
24 */
25 /*
26 * Copyright 2011 Joyent, Inc. All rights reserved.
27 */

29 #ifndef _SYS_MACHCPUVAR_H
30 #define _SYS_MACHCPUVAR_H

32 #ifdef __cplusplus
33 extern "C" {
34 #endif

36 #include <sys/inttypes.h>
37 #include <sys/x_call.h>
38 #include <sys/tss.h>
39 #include <sys/segments.h>
40 #include <sys/rm_platter.h>
41 #include <sys/avintr.h>
42 #include <sys/pte.h>

44 #ifndef _ASM
45 /*
46 * On a virtualized platform a virtual cpu may not be actually
47 * on a physical cpu, especially in situations where a configuration has
48 * more vcpus than pcpus. This function tells us (if it's able) if the
49 * specified vcpu is currently running on a pcpu. Note if it is not
50 * known or not able to determine, it will return the unknown state.
51 */
52 #define VCPU_STATE_UNKNOWN    0
53 #define VCPU_ON_PCPU         1
54 #define VCPU_NOT_ON_PCPU     2

56 extern int vcpu_on_pcpu(processorid_t);

58 /*
59 * Machine specific fields of the cpu struct
60 * defined in common/sys/cpuvar.h.
61 */

```

```

62 * Note: This is kinda kludgy but seems to be the best
63 * of our alternatives.
64 */
65 typedef void *cpu_pri_lev_t;

66 struct cpuid_info;
67 struct cpu_ucode_info;
68 struct cmi_hdl;

70 /*
71 * A note about the hypervisor affinity bits: a one bit in the affinity mask
72 * means the corresponding event channel is allowed to be serviced
73 * by this cpu.
74 */
75 struct xen_evt_data {
76     ulong_t         pending_sel[PIL_MAX + 1]; /* event array selectors */
77     ulong_t         pending_evts[PIL_MAX + 1][sizeof (ulong_t) * 8];
78     ulong_t         evt_affinity[sizeof (ulong_t) * 8]; /* service on cpu */
79 };

81 struct machcpu {
82     /*
83      * x_call fields - used for interprocessor cross calls
84      */
85     struct xc_msg    *xc_msgbox;
86     struct xc_msg    *xc_free;
87     xc_data_t        xc_data;
88     uint32_t         xc_wait_cnt;
89     volatile uint32_t xc_work_cnt;

91     int              mcpu_nodeid; /* node-id */
92     int              mcpu_pri; /* CPU priority */
94     cpu_pri_lev_t    mcpu_pri_data; /* ptr to machine dependent */
95     /* data for setting priority */
96     /* level */

94     struct hat      *mcpu_current_hat; /* cpu's current hat */

96     struct hat_cpu_info *mcpu_hat_info;

98     volatile ulong_t mcpu_tlb_info;

100 /* i86 hardware table addresses that cannot be shared */

102     user_desc_t      *mcpu_gdt; /* GDT */
103     gate_desc_t      *mcpu_idt; /* current IDT */

105     tss_t            *mcpu_tss; /* TSS */

107     kmutex_t         mcpu_ppaddr_mutex;
108     caddr_t          mcpu_caddr1; /* per cpu CADDR1 */
109     caddr_t          mcpu_caddr2; /* per cpu CADDR2 */
110     uint64_t         mcpu_caddr1pte;
111     uint64_t         mcpu_caddr2pte;

113     struct softint   mcpu_softinfo;
114     uint64_t         pil_high_start[HIGH_LEVELS];
115     uint64_t         intrstat[PIL_MAX + 1][2];

117     struct cpuid_info *mcpu_cpi;

119 #if defined(__amd64)
120     greg_t          mcpu_rtmp_rsp; /* syscall: temporary %rsp stash */
121     greg_t          mcpu_rtmp_r15; /* syscall: temporary %r15 stash */
122 #endif

```

```
124     struct vcpu_info *mcpu_vcpu_info;
125     uint64_t      mcpu_gdtpa; /* hypervisor: GDT physical address */

127     uint16_t mcpu_intr_pending; /* hypervisor: pending intrpt levels */
128     uint16_t mcpu_ec_mbox; /* hypervisor: evtchn_dev mailbox */
129     struct xen_evt_data *mcpu_evt_pend; /* hypervisor: pending events */

131     volatile uint32_t *mcpu_mwait; /* MONITOR/MWAIT buffer */
132     void (*mcpu_idle_cpu)(void); /* idle function */
133     uint16_t mcpu_idle_type; /* CPU next idle type */
134     uint16_t max_cstates; /* supported max cstates */

136     struct cpu_ucode_info *mcpu_ucode_info;

138     void *mcpu_pm_mach_state;
139     struct cmi_hdl *mcpu_cmi_hdl;
140     void *mcpu_mach_ctx_ptr;

142     /*
143     * A stamp that is unique per processor and changes
144     * whenever an interrupt happens. Useful for detecting
145     * if a section of code gets interrupted.
146     * The high order 16 bits will hold the cpu->cpu_id.
147     * The low order bits will be incremented on every interrupt.
148     */
149     volatile uint32_t mcpu_istamp;
150 };

152 #define NINTR_THREADS (LOCK_LEVEL-1) /* number of interrupt threads */
153 #define MWAIT_HALTED (1) /* mcpu_mwait set when halting */
154 #define MWAIT_RUNNING (0) /* mcpu_mwait set to wakeup */
155 #define MWAIT_WAKEUP_IPI (2) /* need IPI to wakeup */
156 #define MWAIT_WAKEUP(cpu) (*(cpu->cpu_m.mcpu_mwait) = MWAIT_RUNNING)

158 #endif /* _ASM */

160 /* Please DON'T add any more of this namespace-poisoning sewage here */

162 #define cpu_nodeid cpu_m.mcpu_nodeid
163 #define cpu_pri cpu_m.mcpu_pri
164 #define cpu_current_hat cpu_m.mcpu_current_hat
165 #define cpu_hat_info cpu_m.mcpu_hat_info
166 #define cpu_ppaddr_mutex cpu_m.mcpu_ppaddr_mutex
167 #define cpu_gdt cpu_m.mcpu_gdt
168 #define cpu_idt cpu_m.mcpu_idt
169 #define cpu_tss cpu_m.mcpu_tss
170 #define cpu_ldt cpu_m.mcpu_ldt
171 #define cpu_caddr1 cpu_m.mcpu_caddr1
172 #define cpu_caddr2 cpu_m.mcpu_caddr2
173 #define cpu_softinfo cpu_m.mcpu_softinfo
174 #define cpu_caddr1pte cpu_m.mcpu_caddr1pte
175 #define cpu_caddr2pte cpu_m.mcpu_caddr2pte

177 #ifdef __cplusplus
178 }
_____unchanged_portion_omitted_____
```