

new/usr/src/cmd/mdb/common/kmdb/kaif\_start.c

1

```
*****
12924 Wed Aug 15 14:55:58 2018
new/usr/src/cmd/mdb/common/kmdb/kaif_start.c
9736 kmdb tortures via single-step miscellaneous trap
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 2007 Sun Microsystems, Inc. All rights reserved.
23 * Use is subject to license terms.
24 * Copyright 2018 Joyent, Inc.
25 */
26 #pragma ident "%Z%M% %I% %E% SMI"
27 /*
28 * The main CPU-control loops, used to control masters and slaves.
29 */
31 #include <sys/types.h>
33 #include <kmdb/kaif.h>
34 #include <kmdb/kaif_start.h>
35 #include <kmdb/kmdb_asmutil.h>
36 #include <kmdb/kmdb_dpi_impl.h>
37 #include <kmdb/kmdb_kdi.h>
39 #define KAIF_SLAVE_CMD_SPIN 0
40 #define KAIF_SLAVE_CMD_SWITCH 1
41 #define KAIF_SLAVE_CMD_RESUME 2
42 #define KAIF_SLAVE_CMD_FLUSH 3
43 #define KAIF_SLAVE_CMD_REBOOT 4
44 #if defined(__sparc)
45 #define KAIF_SLAVE_CMD_ACK 5
46 #endif
49 /*
50 * Used to synchronize attempts to set kaif_master_cpuid. kaif_master_cpuid may
51 * be read without kaif_master_lock, and may be written by the current master
52 * CPU.
53 */
54 int kaif_master_cpuid = KAIF_MASTER_CPUID_UNSET;
55 static uintptr_t kaif_master_lock = 0;
57 /*
```

new/usr/src/cmd/mdb/common/kmdb/kaif\_start.c

2

```
58 * Used to ensure that all CPUs leave the debugger together. kaif_loop_lock must
59 * be held to write kaif_looping, but need not be held to read it.
60 */
61 static volatile uint_t kaif_looping;
62 static uintptr_t kaif_loop_lock;
64 static volatile int kaif_slave_cmd;
65 static volatile int kaif_slave_tgt; /* target cpuid for CMD_SWITCH */
67 static void
68 kaif_lock_enter(uintptr_t *lock)
69 {
70     while (cas(lock, 0, 1) != 0)
71         continue;
72     membar_producer();
73 }
74
75 unchanged_portion_omitted
288 int
289 kaif_main_loop(kaif_cpusave_t *cpusave)
290 {
291     int cmd;
293     if (kaif_master_cpuid == KAIF_MASTER_CPUID_UNSET) {
295         /*
296          * Special case: Unload requested before first debugger entry.
297          * Don't stop the world, as there's nothing to clean up that
298          * can't be handled by the running kernel.
299          */
300         if (!kmdb_dpi_resume_requested &&
301             kmdb_kdi_get_unload_request()) {
302             cpusave->krs_cpu_state = KAIF_CPU_STATE_NONE;
303             return (KAIF_CPU_CMD_RESUME);
304         }
306         /*
307          * We're a slave with no master, so just resume. This can
308          * happen if, prior to this, two CPUs both raced through
309          * kdi_cmint() - for example, a breakpoint on a frequently
310          * called function. The loser will be redirected to the slave
311          * loop; note that the event itself is lost at this point.
312          *
313          * The winner will then cross-call that slave, but it won't
314          * actually be received until the slave returns to the kernel
315          * and enables interrupts. We'll then come back in via
316          * kdi_slave_entry() and hit this path.
317          * Special case: Unload requested before first debugger
318          * entry. Don't stop the world, as there's nothing to
319          * clean up that can't be handled by the running kernel.
320          */
321         if (cpusave->krs_cpu_state == KAIF_CPU_STATE_SLAVE) {
322             cpusave->krs_cpu_state = KAIF_CPU_STATE_NONE;
323             return (KAIF_CPU_CMD_RESUME);
324         }
325     }
326     kaif_select_master(cpusave);
327
328 #ifdef __sparc
329     if (kaif_master_cpuid == cpusave->krs_cpu_id) {
330         /*
331          * Everyone has arrived, so we can disarm the post-PROM
332          * entry point.
333          */
334         *kaif_promexitarmp = 0;
335         membar_producer();
336     }
337 #endif
338 }
```

```

333     }
334 #endif
335 } else if (kaif_master_cpuid == cpusave->krs_cpu_id) {
336     cpusave->krs_cpu_state = KAIF_CPU_STATE_MASTER;
337 } else {
338     cpusave->krs_cpu_state = KAIF_CPU_STATE_SLAVE;
339 }
341 cpusave->krs_cpu_flushed = 0;
343 kaif_lock_enter(&kaif_loop_lock);
344 kaif_looping++;
345 kaif_lock_exit(&kaif_loop_lock);
347 /*
348  * We know who the master and slaves are, so now they can go off
349  * to their respective loops.
350  */
351 do {
352     if (kaif_master_cpuid == cpusave->krs_cpu_id)
353         cmd = kaif_master_loop(cpusave);
354     else
355         cmd = kaif_slave_loop(cpusave);
356 } while (cmd == KAIF_CPU_CMD_SWITCH);
358 kaif_lock_enter(&kaif_loop_lock);
359 kaif_looping--;
360 kaif_lock_exit(&kaif_loop_lock);
362 cpusave->krs_cpu_state = KAIF_CPU_STATE_NONE;
364 if (cmd == KAIF_CPU_CMD_RESUME) {
365     /*
366      * By this point, the master has directed the slaves to resume,
367      * and everyone is making their way to this point. We're going
368      * to block here until all CPUs leave the master and slave
369      * loops. When all have arrived, we'll turn them all loose.
370      * This barrier is required for two reasons:
371      *
372      * 1. There exists a race condition whereby a CPU could reenter
373      *    the debugger while another CPU is still in the slave loop
374      *    from this debugger entry. This usually happens when the
375      *    current master releases the slaves, and makes it back to
376      *    the world before the slaves notice the release. The
377      *    former master then triggers a debugger entry, and attempts
378      *    to stop the slaves for this entry before they've even
379      *    resumed from the last one. When the slaves arrive here,
380      *    they'll have re-disabled interrupts, and will thus ignore
381      *    cross-calls until they finish resuming.
382      *
383      * 2. At the time of this writing, there exists a SPARC bug that
384      *    causes an apparently unsolicited interrupt vector trap
385      *    from OBP to one of the slaves. This wouldn't normally be
386      *    a problem but for the fact that the cross-called CPU
387      *    encounters some sort of failure while in OBP. OBP
388      *    recovers by executing the debugger-hook word, which sends
389      *    the slave back into the debugger, triggering a debugger
390      *    fault. This problem seems to only happen during resume,
391      *    the result being that all CPUs save for the cross-called
392      *    one make it back into the world, while the cross-called
393      *    one is stuck at the debugger fault prompt. Leave the
394      *    world in that state too long, and you'll get a mondo
395      *    timeout panic. If we hold everyone here, we can give the
396      *    the user a chance to trigger a panic for further analysis.
397      *    To trigger the bug, "pool_unlock:b :c" and "while : ; do
398      *    prsset -p ; done".

```

```

399     *
400     * When the second item is fixed, the barrier can move into
401     * kaif_select_master(), immediately prior to the setting of
402     * kaif_master_cpuid.
403     */
404     while (kaif_looping != 0)
405         continue;
406 }
408     return (cmd);
409 }

```

unchanged portion omitted

new/usr/src/uts/intel/amd64/sys/kdi\_regs.h

1

```
*****
2853 Wed Aug 15 14:55:58 2018
new/usr/src/uts/intel/amd64/sys/kdi_regs.h
9736 kmdb tortures via single-step miscellaneous trap
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 */

22 /*
23 * Copyright 2007 Sun Microsystems, Inc. All rights reserved.
24 * Use is subject to license terms.
25 *
26 * Copyright 2018 Joyent, Inc.
27 */

29 #ifndef _AMD64_SYS_KDI_REGS_H
30 #define _AMD64_SYS_KDI_REGS_H

32 #include <sys/stddef.h>

34 #ifdef __cplusplus
35 extern "C" {
36 #endif

38 /*
39 * A modified version of struct regs layout.
40 */

42 #define KDIREG_SAVFP 0
43 #define KDIREG_SAVPC 1
44 #define KDIREG_RDI 2
45 #define KDIREG_RSI 3
46 #define KDIREG_RDX 4
47 #define KDIREG_RCX 5
48 #define KDIREG_R8 6
49 #define KDIREG_R9 7
50 #define KDIREG_RAX 8
51 #define KDIREG_RBX 9
52 #define KDIREG_RBP 10
53 #define KDIREG_R10 11
54 #define KDIREG_R11 12
55 #define KDIREG_R12 13
56 #define KDIREG_R13 14
57 #define KDIREG_R14 15
58 #define KDIREG_R15 16
59 #define KDIREG_FSBASE 17
```

new/usr/src/uts/intel/amd64/sys/kdi\_regs.h

2

```
60 #define KDIREG_GSBASE 18
61 #define KDIREG_KGSBASE 19
62 #define KDIREG_CR2 20
63 #define KDIREG_CR3 21
64 #define KDIREG_DS 22
65 #define KDIREG_ES 23
66 #define KDIREG_FS 24
67 #define KDIREG_GS 25
68 #define KDIREG_TRAPNO 26
69 #define KDIREG_ERR 27
70 #define KDIREG_RIP 28
71 #define KDIREG_CS 29
72 #define KDIREG_RFLAGS 30
73 #define KDIREG_RSP 31
74 #define KDIREG_SS 32

76 #define KDIREG_NGREG (KDIREG_SS + 1)

78 #define KDIREG_PC KDIREG_RIP
79 #define KDIREG_SP KDIREG_RSP
80 #define KDIREG_FP KDIREG_RBP

82 #if !defined(_ASM)

84 /*
85  * Handy for debugging krs_gregs; keep in sync with the KDIREG_* above.
86  */
87 typedef struct {
88     greg_t kr_savfp;
89     greg_t kr_savpc;
90     greg_t kr_rdi;
91     greg_t kr_rsi;
92     greg_t kr_rdx;
93     greg_t kr_rcx;
94     greg_t kr_r8;
95     greg_t kr_r9;
96     greg_t kr_rax;
97     greg_t kr_rbx;
98     greg_t kr_rbp;
99     greg_t r_r10;
100    greg_t r_r11;
101    greg_t r_r12;
102    greg_t r_r13;
103    greg_t r_r14;
104    greg_t r_r15;
105    greg_t kr_fsbases;
106    greg_t kr_gsbases;
107    greg_t kr_kgsbases;
108    greg_t kr_cr2;
109    greg_t kr_cr3;
110    greg_t kr_ds;
111    greg_t kr_es;
112    greg_t kr_fs;
113    greg_t kr_gs;
114    greg_t kr_trapno;
115    greg_t kr_err;
116    greg_t kr_rip;
117    greg_t kr_cs;
118    greg_t kr_rflags;
119    greg_t kr_rsp;
120    greg_t kr_ss;
121 } kdiregs_t;

123 #if defined(KERNEL)
124 CTASSERT(offsetof(kdiregs_t, kr_ss) == ((KDIREG_NGREG - 1) * sizeof(greg_t)));
125 #endif
```

```
127 #endif /* !_ASM */  
129 #ifdef __cplusplus  
130 }  
_____unchanged_portion_omitted_____
```

```

*****
18997 Wed Aug 15 14:55:59 2018
new/usr/src/uts/intel/kdi/kdi_asm.s
9736 kmdb tortures via single-step miscellaneous trap
Reviewed by: Robert Mustacchi <rm@joyent.com>
Reviewed by: Jerry Jelinek <jerry.jelinek@joyent.com>
*****
unchanged_portion_omitted

```

```

370 /*
371 * The cross-call handler for slave CPUs.
372 *
373 * The debugger is single-threaded, so only one CPU, called the master, may be
374 * running it at any given time. The other CPUs, known as slaves, spin in a
375 * busy loop until there's something for them to do. This is the entry point
376 * for the slaves - they'll be sent here in response to a cross-call sent by the
377 * master.
378 */

380 ENTRY_NP(kdi_slave_entry)

382 /*
383 * Cross calls are implemented as function calls, so our stack currently
384 * looks like one you'd get from a zero-argument function call. That
385 * is, there's the return %rip at %rsp, and that's about it. We need
386 * to make it look like an interrupt stack. When we first save, we'll
387 * reverse the saved %ss and %rip, which we'll fix back up when we've
388 * freed up some general-purpose registers. We'll also need to fix up
389 * the saved %rsp.
390 */

392 pushq %rsp          /* pushed value off by 8 */
393 pushfq
394 CLI(%rax)
395 pushq $KCS_SEL
396 clrq %rax
397 movw %ss, %ax
398 pushq %rax          /* rip should be here */
399 pushq $-1          /* phony trap error code */
400 pushq $-1          /* phony trap number */

402 subq $REG_OFF(KDIREG_TRAPNO), %rsp
403 KDI_SAVE_REGS(%rsp)

405 movq %cr3, %rax
406 movq %rax, REG_OFF(KDIREG_CR3)(%rsp)

408 movq REG_OFF(KDIREG_SS)(%rsp), %rax
409 movq %rax, REG_OFF(KDIREG_SAVPC)(%rsp)
410 xchgq REG_OFF(KDIREG_RIP)(%rsp), %rax
411 movq %rax, REG_OFF(KDIREG_SS)(%rsp)

413 movq REG_OFF(KDIREG_RSP)(%rsp), %rax
414 addq $8, %rax
415 movq %rax, REG_OFF(KDIREG_RSP)(%rsp)

417 /*
418 * We've saved all of the general-purpose registers, and have a stack
419 * that is irettable (after we strip down to the error code)
420 */

422 GET_CPUSAVE_ADDR      /* %rax = cpusave, %rbx = CPU ID */

424 ADVANCE_CRUMB_POINTER(%rax, %rcx, %rdx)

426 ADD_CRUMB(%rax, KRM_CPU_STATE, $KDI_CPU_STATE_SLAVE, %rdx)

```

```

428 movq REG_OFF(KDIREG_RIP)(%rsp), %rcx
429 ADD_CRUMB(%rax, KRM_PC, %rcx, %rdx)
430 movq REG_OFF(KDIREG_RSP)(%rsp), %rcx
431 ADD_CRUMB(%rax, KRM_SP, %rcx, %rdx)
432 ADD_CRUMB(%rax, KRM_TRAPNO, $-1, %rdx)

434 movq $KDI_CPU_STATE_SLAVE, KRS_CPU_STATE(%rax)

436 pushq %rax
437 jmp kdi_save_common_state

439 SET_SIZE(kdi_slave_entry)
unchanged_portion_omitted

```

```

*****
12331 Wed Aug 15 14:55:59 2018
new/usr/src/uts/intel/kdi/kdi_idt.c
9736 kmdb tortures via single-step miscellaneous trap
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 2009 Sun Microsystems, Inc. All rights reserved.
23 * Use is subject to license terms.
24 *
25 * Copyright 2018 Joyent, Inc.
26 */
27
28 /*
29 * Management of KMDB's IDT, which is installed upon KMDB activation.
30 *
31 * Debugger activation has two flavors, which cover the cases where KMDB is
32 * loaded at boot, and when it is loaded after boot. In brief, in both cases,
33 * the KDI needs to interpose upon several handlers in the IDT. When
34 * mod-loaded KMDB is deactivated, we undo the IDT interposition, restoring the
35 * handlers to what they were before we started.
36 *
37 * We also take over the entirety of IDT (except the double-fault handler) on
38 * the active CPU when we're in kmdb so we can handle things like page faults
39 * sensibly.
40 *
41 * Boot-loaded KMDB
42 *
43 * When we're first activated, we're running on boot's IDT. We need to be able
44 * to function in this world, so we'll install our handlers into boot's IDT.
45 * This is a little complicated: we're using the fake cpu_t set up by
46 * boot_kdi_tmpinit(), so we can't access cpu_idt directly. Instead,
47 * kdi_idt_write() notices that cpu_idt is NULL, and works around this problem.
48 *
49 * Later, when we're about to switch to the kernel's IDT, it'll call us via
50 * kdi_idt_sync(), allowing us to add our handlers to the new IDT. While
51 * boot-loaded KMDB can't be unloaded, we still need to save the descriptors we
52 * replace so we can pass traps back to the kernel as necessary.
53 *
54 * The last phase of boot-loaded KMDB activation occurs at non-boot CPU
55 * startup. We will be called on each non-boot CPU, thus allowing us to set up
56 * any watchpoints that may have been configured on the boot CPU and interpose
57 * on the given CPU's IDT. We don't save the interposed descriptors in this
58 * case -- see kdi_cpu_init() for details.
59 *

```

```

60 * Mod-loaded KMDB
61 *
62 * This style of activation is much simpler, as the CPUs are already running,
63 * and are using their own copy of the kernel's IDT. We simply interpose upon
64 * each CPU's IDT. We save the handlers we replace, both for deactivation and
65 * for passing traps back to the kernel. Note that for the hypervisors'
66 * benefit, we need to xcall to the other CPUs to do this, since we need to
67 * actively set the trap entries in its virtual IDT from that vcpu's context
68 * rather than just modifying the IDT table from the CPU running kdi_activate().
69 */
70
71 #include <sys/types.h>
72 #include <sys/segments.h>
73 #include <sys/trap.h>
74 #include <sys/cpuvar.h>
75 #include <sys/reboot.h>
76 #include <sys/sunddi.h>
77 #include <sys/archsystem.h>
78 #include <sys/kdi_impl.h>
79 #include <sys/x_call.h>
80 #include <ia32/sys/psw.h>
81 #include <vm/hat_i86.h>
82
83 #define KDI_GATE_NVECS 3
84
85 #define KDI_IDT_NOSAVE 0
86 #define KDI_IDT_SAVE 1
87
88 #define KDI_IDT_DTYPE_KERNEL 0
89 #define KDI_IDT_DTYPE_BOOT 1
90
91 /* Solely to keep kdiregs_t in the CTF, otherwise unused. */
92 kdiregs_t kdi_regs;
93
94 kdi_cpusave_t *kdi_cpusave;
95 int kdi_ncpusave;
96
97 static kdi_main_t kdi_kmdb_main;
98
99 kdi_drreg_t kdi_drreg;
100
101 #ifndef __amd64
102 /* Used to track the current set of valid kernel selectors. */
103 uint32_t kdi_cs;
104 uint32_t kdi_ds;
105 uint32_t kdi_fs;
106 uint32_t kdi_gs;
107 #endif
108
109 uintptr_t kdi_kernel_handler;
110
111 int kdi_trap_switch;
112
113 #define KDI_MEMRANGES_MAX 2
114
115 kdi_memrange_t kdi_memranges[KDI_MEMRANGES_MAX];
116 int kdi_nmemranges;
117
118 typedef void idt_hdlr_f(void);
119
120 extern idt_hdlr_f kdi_trap0, kdi_trap1, kdi_int2, kdi_trap3, kdi_trap4;
121 extern idt_hdlr_f kdi_trap5, kdi_trap6, kdi_trap7, kdi_trap9;
122 extern idt_hdlr_f kdi_traperr10, kdi_traperr11, kdi_traperr12;
123 extern idt_hdlr_f kdi_traperr13, kdi_traperr14, kdi_trap16, kdi_traperr17;
124 extern idt_hdlr_f kdi_trap18, kdi_trap19, kdi_trap20, kdi_ivct32;
125 extern idt_hdlr_f kdi_invaltrap;

```

new/usr/src/uts/intel/kdi/kdi\_idt.c

3

```
126 extern size_t kdi_ivct_size;

128 typedef struct kdi_gate_spec {
129     uint_t kgs_vec;
130     uint_t kgs_dpl;
131 } kdi_gate_spec_t;
unchanged_portion_omitted
```