

```

*****
24136 Mon Jul 30 12:50:08 2018
new/usr/src/uts/i86pc/ml/kpti_trampoline.s
9685 KPTI %cr3 handling needs fixes
*****
_____unchanged_portion_omitted_____

165 #define SET_KERNEL_CR3(spillreg) \
166     mov     %cr3, spillreg; \
167     mov     spillreg, %gs:CPU_KPTI_TR_CR3; \
168     mov     %gs:CPU_KPTI_UCR3, spillreg; \
169     cmp     $0, spillreg; \
170     je     2f; \
171     mov     spillreg, %cr3; \
172 2:

174 #if DEBUG
175 #define SET_USER_CR3(spillreg) \
176     mov     %cr3, spillreg; \
177     mov     spillreg, %gs:CPU_KPTI_TR_CR3; \
178     mov     %gs:CPU_KPTI_UCR3, spillreg; \
179     mov     spillreg, %cr3 \
180 #else
181 #define SET_USER_CR3(spillreg) \
182     mov     %gs:CPU_KPTI_UCR3, spillreg; \
183     mov     spillreg, %cr3 \
184 #endif

186 #define PIVOT_KPTI_STK(spillreg) \
187     mov     %rsp, spillreg; \
188     mov     %gs:CPU_KPTI_RET_RSP, %rsp; \
189     pushq   T_FRAMERET_SS(spillreg); \
190     pushq   T_FRAMERET_RSP(spillreg); \
191     pushq   T_FRAMERET_RFLAGS(spillreg); \
192     pushq   T_FRAMERET_CS(spillreg); \
193     pushq   T_FRAMERET_RIP(spillreg)

196 #define INTERRUPT_TRAMPOLINE_P(errpush) \
197     pushq   %r13; \
198     pushq   %r14; \
199     subq   $KPTI_R14, %rsp; \
200     /* Save current %cr3. */ \
201     mov     %cr3, %r14; \
202     mov     %r14, KPTI_TR_CR3(%rsp); \
203     \
204     cmpw   $KCS_SEL, KPTI_CS(%rsp); \
205     je     3f; \
206 1: \
207     /* Change to the "kernel" %cr3 */ \
208     mov     KPTI_KCR3(%rsp), %r14; \
209     cmp     $0, %r14; \
210     je     2f; \
211     mov     %r14, %cr3; \
212 2: \
213     /* Get our cpu_t in %r13 */ \
214     mov     %rsp, %r13; \
215     and     $(~(MMU_PAGESIZE - 1)), %r13; \
216     subq   $CPU_KPTI_START, %r13; \
217     /* Use top of the kthread stk */ \
218     mov     CPU_THREAD(%r13), %r14; \
219     mov     T_STACK(%r14), %r14; \
220     addq   $REGSIZE+MINFRAME, %r14; \
221     jmp     4f; \
222 3: \
223     /* Check the %rsp in the frame. */ \

```

```

224     /* Is it above kernel base? */ \
225     mov     kpti_kbase, %r14; \
226     cmp     %r14, KPTI_RSP(%rsp); \
227     jb     1b; \
228     /* Use the %rsp from the trap frame */ \
229     mov     KPTI_RSP(%rsp), %r14; \
230     and     $(~0xf), %r14; \
231 4: \
232     mov     %rsp, %r13; \
233     /* %r14 contains our destination stk */ \
234     mov     %r14, %rsp; \
235     pushq   KPTI_SS(%r13); \
236     pushq   KPTI_RSP(%r13); \
237     pushq   KPTI_RFLAGS(%r13); \
238     pushq   KPTI_CS(%r13); \
239     pushq   KPTI_RIP(%r13); \
240     errpush; \
241     mov     KPTI_R14(%r13), %r14; \
242     mov     KPTI_R13(%r13), %r13

244 #define INTERRUPT_TRAMPOLINE_NOERR \
245     INTERRUPT_TRAMPOLINE_P(/**/)

247 #define INTERRUPT_TRAMPOLINE \
248     INTERRUPT_TRAMPOLINE_P(pushq KPTI_ERR(%r13))

250 /*
251 * This is used for all interrupts that can plausibly be taken inside another
252 * interrupt and are using a kpti_frame stack (so #BP, #DB, #GP, #PF, #SS).
253 *
254 * We also use this for #NP, even though it uses the standard IST: the
255 * additional %rsp checks below will catch when we get an exception doing an
256 * iret to userspace with a bad %cs/%ss. This appears as a kernel trap, and
257 * only later gets redirected via kern_gpfault().
258 *
259 * We check for whether we took the interrupt while in another trampoline, in
260 * which case we need to use the kthread stack.
261 */
262 #define DBG_INTERRUPT_TRAMPOLINE_P(errpush) \
263     pushq   %r13; \
264     pushq   %r14; \
265     subq   $KPTI_R14, %rsp; \
266     /* Check for clobbering */ \
267     cmp     $0, KPTI_FLAG(%rsp); \
268     je     1f; \
269     /* Don't worry, this totally works */ \
270     int     $8; \
271 1: \
272     movq   $1, KPTI_FLAG(%rsp); \
273     /* Save current %cr3. */ \
274     mov     %cr3, %r14; \
275     mov     %r14, KPTI_TR_CR3(%rsp); \
276     \
277     cmpw   $KCS_SEL, KPTI_CS(%rsp); \
278     je     4f; \
279 2: \
280     /* Change to the "kernel" %cr3 */ \
281     mov     KPTI_KCR3(%rsp), %r14; \
282     cmp     $0, %r14; \
283     je     3f; \
284     mov     %r14, %cr3; \
285 3: \
286     /* Get our cpu_t in %r13 */ \
287     mov     %rsp, %r13; \
288     and     $(~(MMU_PAGESIZE - 1)), %r13; \
289     subq   $CPU_KPTI_START, %r13; \

```

```

290      /* Use top of the kthread stk */
291      mov     CPU_THREAD(%r13), %r14;
292      mov     T_STACK(%r14), %r14;
293      addq   $REGSIZE+MINFRAME, %r14;
294      jmp    6f;
295 4:
296      /* Check the %rsp in the frame. */
297      /* Is it above kernel base? */
298      /* If not, treat as user. */
299      mov     kpti_kbase, %r14;
300      cmp     %r14, KPTI_RSP(%rsp);
301      jb     2b;
302      /* Is it within the kpti_frame page? */
303      /* If it is, treat as user interrupt */
304      mov     %rsp, %r13;
305      and     $(~(MMU_PAGESIZE - 1)), %r13;
306      mov     KPTI_RSP(%rsp), %r14;
307      and     $(~(MMU_PAGESIZE - 1)), %r14;
308      cmp     %r13, %r14;
309      je     2b;
310      /* Were we in trampoline code? */
311      leaq   kpti_tramp_start, %r14;
312      cmp     %r14, KPTI_RIP(%rsp);
313      jb     5f;
314      leaq   kpti_tramp_end, %r14;
315      cmp     %r14, KPTI_RIP(%rsp);
316      ja     5f;
317      /* If we were, change %cr3: we might */
318      /* have interrupted before it did. */
319      mov     KPTI_KCR3(%rsp), %r14;
320      mov     %r14, %cr3;
321 5:
322      /* Use the %rsp from the trap frame */
323      mov     KPTI_RSP(%rsp), %r14;
324      and     $(~0xf), %r14;
325 6:
326      mov     %rsp, %r13;
327      /* %r14 contains our destination stk */
328      mov     %r14, %rsp;
329      pushq  KPTI_SS(%r13);
330      pushq  KPTI_RSP(%r13);
331      pushq  KPTI_RFLAGS(%r13);
332      pushq  KPTI_CS(%r13);
333      pushq  KPTI_RIP(%r13);
334      errpush;
335      mov     KPTI_R14(%r13), %r14;
336      movq   $0, KPTI_FLAG(%r13);
337      mov     KPTI_R13(%r13), %r13

339 #define DBG_INTERRUPT_TRAMPOLINE_NOERR
340     DBG_INTERRUPT_TRAMPOLINE_P(**/)

342 #define DBG_INTERRUPT_TRAMPOLINE
343     DBG_INTERRUPT_TRAMPOLINE_P(pushq KPTI_ERR(%r13))

345      /*
346      * These labels (_start and _end) are used by trap.c to determine if
347      * we took an interrupt like an NMI during the return process.
348      */
349      .global tr_sysc_ret_start
350      tr_sysc_ret_start:

352      /*
353      * Syscall return trampolines.
354      *
355      * These are expected to be called on the kernel %gs. tr_sysret[ql] are

```

```

356      * called after %rsp is changed back to the user value, so we have no
357      * stack to work with. tr_sysexit has a kernel stack (but has to
358      * preserve rflags, soooo).
359      */
360      ENTRY_NP(tr_sysretq)
361      cmpq   $1, kpti_enable
362      jne    1f

364      mov     %r13, %gs:CPU_KPTI_R13
365      SET_USER_CR3(%r13)
366      mov     %gs:CPU_KPTI_R13, %r13
367      /* Zero these to make sure they didn't leak from a kernel trap */
368      movq   $0, %gs:CPU_KPTI_R13
369      movq   $0, %gs:CPU_KPTI_R14
370 1:
371      swapgs
372      sysretq
373      SET_SIZE(tr_sysretq)
unchanged_portion_omitted

649      MK_INTR_TRAMPOLINE_NOERR(div0trap)
650      MK_DBG_INTR_TRAMPOLINE_NOERR(dbgtrap)
651      MK_DBG_INTR_TRAMPOLINE_NOERR(brktrap)
652      MK_INTR_TRAMPOLINE_NOERR(ovflotrap)
653      MK_INTR_TRAMPOLINE_NOERR(boundstrap)
654      MK_INTR_TRAMPOLINE_NOERR(invoptrap)
655      MK_INTR_TRAMPOLINE_NOERR(ndptrap)
656      MK_INTR_TRAMPOLINE(invtsstrap)
657      MK_DBG_INTR_TRAMPOLINE(segnptrap)
658      MK_INTR_TRAMPOLINE(segnptrap)
659      MK_DBG_INTR_TRAMPOLINE(stktrap)
660      MK_DBG_INTR_TRAMPOLINE(gptrap)
661      MK_INTR_TRAMPOLINE_NOERR(resvtrap)
662      MK_INTR_TRAMPOLINE_NOERR(ndperr)
663      MK_INTR_TRAMPOLINE(achktrap)
664      MK_INTR_TRAMPOLINE_NOERR(xmtrap)
665      MK_INTR_TRAMPOLINE_NOERR(invaltrap)
666      MK_INTR_TRAMPOLINE_NOERR(fasttrap)
667      MK_INTR_TRAMPOLINE_NOERR(dtrace_ret)

669      /*
670      * These are special because they can interrupt other traps, and
671      * each other. We don't need to pivot their stacks, because they have
672      * dedicated IST stack space, but we need to change %cr3.
673      */
674      ENTRY_NP(tr_nmiint)
675      pushq  %r13
676      mov     kpti_safe_cr3, %r13
677      mov     %r13, %cr3
678      popq   %r13
679      jmp    nmiint
680      SET_SIZE(tr_nmiint)
unchanged_portion_omitted

```

```

*****
50081 Mon Jul 30 12:50:09 2018
new/usr/src/uts/i86pc/ml/locore.s
9685 KPTI %cr3 handling needs fixes
*****
_____unchanged_portion_omitted_____

1080 #endif /* __lint */
1081 #endif /* !_amd64 */

1084 /*
1085  * For stack layout, see privregs.h
1086  * When cmntrap gets called, the error code and trap number have been pushed.
1087  * When cmntrap_pushed gets called, the entire struct regs has been pushed.
1088  */

1090 #if defined(__lint)

1092 /* ARGSUSED */
1093 void
1094 cmntrap()
1095 {}

1097 #else /* __lint */

1099     .globl trap                /* C handler called below */

1101 #if defined(__amd64)

1103     ENTRY_NP2(cmntrap, _cmntrap)

1105     INTR_PUSH

1107     ALTENTRY(cmntrap_pushed)

1109     movq    %rsp, %rbp

1111     /*
1112     * - if this is a #pf i.e. T_PGFLT, %r15 is live
1113     *   and contains the faulting address i.e. a copy of %cr2
1114     * - if this is a #db i.e. T_SGLSTP, %r15 is live
1115     *   and contains the value of %db6
1116     */

1119     TRACE_PTR(%rdi, %rbx, %ebx, %rcx, $TT_TRAP) /* Uses labels 8 and 9 */
1120     TRACE_REGS(%rdi, %rsp, %rbx, %rcx) /* Uses label 9 */
1121     TRACE_STAMP(%rdi) /* Clobbers %eax, %edx, uses 9 */

1123     /*
1124     * We must first check if DTrace has set its NOFAULT bit. This
1125     * regrettably must happen before the trap stack is recorded, because
1126     * this requires a call to getpcstack() and may induce recursion if an
1127     * fbt::getpcstack: enabling is inducing the bad load.
1128     */
1129     movl    %gs:CPU_ID, %eax
1130     shlq   $CPU_CORE_SHIFT, %rax
1131     leaq   cpu_core(%rip), %r8
1132     addq   %r8, %rax
1133     movw   CPUC_DTRACE_FLAGS(%rax), %cx
1134     testw  $CPU_DTRACE_NOFAULT, %cx
1135     jnz    .dtrace_induced

1137     TRACE_STACK(%rdi)

```

```

1139     movq   %rbp, %rdi
1140     movq   %r15, %rsi
1141     movl   %gs:CPU_ID, %edx

1143     /*
1144     * We know that this isn't a DTrace non-faulting load; we can now safely
1145     * reenale interrupts. (In the case of pagefaults, we enter through an
1146     * interrupt gate.)
1147     */
1148     ENABLE_INTR_FLAGS

1150     call   trap                /* trap(rp, addr, cpuid) handles all traps */
1151     jmp    _sys_rtt

1153 .dtrace_induced:
1154     cmpw   $KCS_SEL, REGOFF_CS(%rbp) /* test CS for user-mode trap */
1155     jne    3f                  /* if from user, panic */

1157     cmpl   $T_PGFLT, REGOFF_TRAPNO(%rbp)
1158     je     1f

1160     cmpl   $T_GPFLT, REGOFF_TRAPNO(%rbp)
1161     je     0f

1163     cmpl   $T_ILLINST, REGOFF_TRAPNO(%rbp)
1164     je     0f

1166     cmpl   $T_ZERODIV, REGOFF_TRAPNO(%rbp)
1167     jne    4f                  /* if not PF/GP/UD/DE, panic */

1169     orw   $CPU_DTRACE_DIVZERO, %cx
1170     movw  %cx, CPUC_DTRACE_FLAGS(%rax)
1171     jmp   2f

1173     /*
1174     * If we've taken a GPF, we don't (unfortunately) have the address that
1175     * induced the fault. So instead of setting the fault to BADADDR,
1176     * we'll set the fault to ILL0P.
1177     */
1178 0:
1179     orw   $CPU_DTRACE_ILLOP, %cx
1180     movw  %cx, CPUC_DTRACE_FLAGS(%rax)
1181     jmp   2f

1182 1:
1183     orw   $CPU_DTRACE_BADADDR, %cx
1184     movw  %cx, CPUC_DTRACE_FLAGS(%rax) /* set fault to bad addr */
1185     movq  %r15, CPUC_DTRACE_ILLOP(%rax) /* fault addr is illegal value */
1186
1187 2:
1188     movq   REGOFF_RIP(%rbp), %rdi
1189     movq   %rdi, %r12
1190     call   dtrace_instr_size
1191     addq   %rax, %r12
1192     movq   %r12, REGOFF_RIP(%rbp)
1193     INTR_POP
1194     jmp    tr_iret_auto
1195     /*NOTREACHED*/

1196 3:
1197     leaq   dtrace_badflags(%rip), %rdi
1198     xorl   %eax, %eax
1199     call   panic

1200 4:
1201     leaq   dtrace_badtrap(%rip), %rdi
1202     xorl   %eax, %eax
1203     call   panic
1204     SET_SIZE(cmntrap_pushed)

```

new/usr/src/uts/i86pc/ml/locore.s

3

1205 SET_SIZE(cmnttrap)
 unchanged_portion_omitted

```

*****
18790 Mon Jul 30 12:50:10 2018
new/usr/src/uts/intel/kdi/kdi_asm.s
9685 KPTI %cr3 handling needs fixes
*****
_____unchanged_portion_omitted_____

435 /*
436 * The state of the world:
437 *
438 * The stack has a complete set of saved registers and segment
439 * selectors, arranged in the kdi_regs.h order. It also has a pointer
440 * to our cpusave area.
441 *
442 * We need to save, into the cpusave area, a pointer to these saved
443 * registers. First we check whether we should jump straight back to
444 * the kernel. If not, we save a few more registers, ready the
445 * machine for debugger entry, and enter the debugger.
446 */

448     ENTRY_NP(kdi_save_common_state)

450     popq    %rdi          /* the cpusave area */
451     movq    %rsp, KRS_GREGS(%rdi) /* save ptr to current saved regs */

453     pushq   %rdi
454     call    kdi_trap_pass
455     testq   %rax, %rax
456     jnz     kdi_pass_to_kernel
455     cmpq   $1, %rax
456     je     kdi_pass_to_kernel
457     popq   %rax /* cpusave in %rax */

459     SAVE_IDTGDTR

461 #if !defined(__xpv)
462     /* Save off %cr0, and clear write protect */
463     movq    %cr0, %rcx
464     movq    %rcx, KRS_CR0(%rax)
465     andq    $_BITNOT(CR0_WP), %rcx
466     movq    %rcx, %cr0
467 #endif

469     /* Save the debug registers and disable any active watchpoints */

471     movq    %rax, %r15          /* save cpusave area ptr */
472     movl    $7, %edi
473     call    kdi_dreg_get
474     movq    %rax, KRS_DRCTL(%r15)

476     andq    $_BITNOT(KDIREG_DRCTL_WPALLEN_MASK), %rax
477     movq    %rax, %rsi
478     movl    $7, %edi
479     call    kdi_dreg_set

481     movl    $6, %edi
482     call    kdi_dreg_get
483     movq    %rax, KRS_DRSTAT(%r15)

485     movl    $0, %edi
486     call    kdi_dreg_get
487     movq    %rax, KRS_DROFF(0)(%r15)

489     movl    $1, %edi
490     call    kdi_dreg_get
491     movq    %rax, KRS_DROFF(1)(%r15)

```

```

493     movl    $2, %edi
494     call    kdi_dreg_get
495     movq    %rax, KRS_DROFF(2)(%r15)

497     movl    $3, %edi
498     call    kdi_dreg_get
499     movq    %rax, KRS_DROFF(3)(%r15)

501     movq    %r15, %rax          /* restore cpu save area to rax */

503     clrq    %rbp              /* stack traces should end here */

505     pushq   %rax
506     movq    %rax, %rdi          /* cpusave */

508     call    kdi_debugger_entry

510     /* Pass cpusave to kdi_resume */
511     popq    %rdi

513     jmp     kdi_resume

515     SET_SIZE(kdi_save_common_state)
_____unchanged_portion_omitted_____

572     ENTRY_NP(kdi_pass_to_kernel)

574     popq    %rdi /* cpusave */

576     movq    $KDI_CPU_STATE_NONE, KRS_CPU_STATE(%rdi)

573     /*
574     * We took a trap that should be handled by the kernel, not KMDB.
575     * Find the trap and vector off the right kernel handler. The trap
576     * handler will expect the stack to be in trap order, with %rip being
577     * the last entry, so we'll need to restore all our regs. On i86xpv
578     * we'll need to compensate for XPV_TRAP_POP.
579     *
580     * We're hard-coding the three cases where KMDB has installed permanent
581     * handlers, since after we KDI_RESTORE_REGS(), we don't have registers
582     * to work with; we can't use a global since other CPUs can easily pass
583     * through here at the same time.
584     *
585     * Note that we handle T_DBGENTR since userspace might have tried it.
586     *
587     * The trap handler will expect the stack to be in trap order, with %rip
588     * being the last entry, so we'll need to restore all our regs. On
589     * i86xpv we'll need to compensate for XPV_TRAP_POP.
590     *
591     * %rax on entry is either 1 or 2, which is from kdi_trap_pass().
592     * kdi_cmntnt stashed the original %cr3 into KDIREG_CR3, then (probably)
593     * switched us to the CPU's kf_kernel_cr3. But we're about to call, for
594     * example:
595     *
596     * dbgtrap->trap()->tr_iret_kernel
597     *
598     * which, unlike, tr_iret_kdi, doesn't restore the original %cr3, so
599     * we'll do so here if needed.
600     *
601     * This isn't just a matter of tidiness: for example, consider:
602     *
603     * hat_switch(oldhat=kas.a_hat, newhat=prochat)
604     * setcr3()
605     * reset_kpti()
606     * *brktrap* due to fbt on reset_kpti:entry

```

```
603      *
604      * Here, we have the new hat's %cr3, but we haven't yet updated
605      * kf_kernel_cr3 (so its currently kas's). So if we don't restore here,
606      * we'll stay on kas's cr3 value on returning from the trap: not good if
607      * we fault on a userspace address.
608      */
609      ENTRY_NP(kdi_pass_to_kernel)

611      popq    %rdi /* cpusave */
612      movq    $KDI_CPU_STATE_NONE, KRS_CPU_STATE(%rdi)
613      movq    KRS_GREGS(%rdi), %rsp

615      cmpq    $2, %rax
616      jne     no_restore_cr3
617      movq    REG_OFF(KDIREG_CR3)(%rsp), %r11
618      movq    %r11, %cr3

620 no_restore_cr3:
621      movq    REG_OFF(KDIREG_TRAPNO)(%rsp), %rdi

623      cmpq    $T_SGLSTP, %rdi
624      je      kdi_pass_dbgtrap
625      je      1f
626      cmpq    $T_BPTFLT, %rdi
627      je      kdi_pass_brktrap
628      je      2f
629      cmpq    $T_DBGENTR, %rdi
630      je      kdi_pass_invaltrap
631      je      3f
632      /*
633      * Hmm, unknown handler.  Somebody forgot to update this when they
634      * added a new trap interposition... try to drop back into kmdb.
635      */
636      int     $T_DBGENTR

637 #define CALL_TRAP_HANDLER(name) \
638     KDI_RESTORE_REGS(%rsp); \
639     /* Discard state, trapno, err */ \
640     addq    $REG_OFF(KDIREG_RIP), %rsp; \
641     XPV_TRAP_PUSH; \
642     jmp     %cs:name

643 kdi_pass_dbgtrap:
644 1:
645     CALL_TRAP_HANDLER(debugtrap)
646     /*NOTREACHED*/

647 kdi_pass_brktrap:
648 2:
649     CALL_TRAP_HANDLER(brktrap)
650     /*NOTREACHED*/

651 kdi_pass_invaltrap:
652 3:
653     CALL_TRAP_HANDLER(invaltrap)
654     /*NOTREACHED*/

655     SET_SIZE(kdi_pass_to_kernel)
656     _____unchanged_portion_omitted_____
```

```
*****
12249 Mon Jul 30 12:50:10 2018
new/usr/src/uts/intel/kdi/kdi_idt.c
9685 KPTI %cr3 handling needs fixes
*****
_____unchanged_portion_omitted_____

368 /*
369 * We receive all breakpoints and single step traps. Some of them, including
370 * those from userland and those induced by DTrace providers, are intended for
371 * the kernel, and must be processed there. We adopt this
372 * ours-until-proven-otherwise position due to the painful consequences of
373 * sending the kernel an unexpected breakpoint or single step. Unless someone
374 * can prove to us that the kernel is prepared to handle the trap, we'll assume
375 * there's a problem and will give the user a chance to debug it.
376 *
377 * If we return 2, then the calling code should restore the trap-time %cr3: that
378 * is, it really is a kernel-originated trap.
379 */
380 int
381 kdi_trap_pass(kdi_cpusave_t *cpusave)
382 {
383     greg_t tt = cpusave->krs_gregs[KDIREG_TRAPNO];
384     greg_t pc = cpusave->krs_gregs[KDIREG_PC];
385     greg_t cs = cpusave->krs_gregs[KDIREG_CS];
386
387     if (USERMODE(cs))
388         return (1);
389
390     if (tt != T_BPTFLT && tt != T_SGLSTP)
391         return (0);
392
393     if (tt == T_BPTFLT && kdi_dtrace_get_state() ==
394         KDI_DTSTATE_DTRACE_ACTIVE)
395         return (2);
396     return (1);
397
398     /*
399     * See the comments in the kernel's T_SGLSTP handler for why we need to
400     * do this.
401     */
402     #if !defined(__xpv)
403     if (tt == T_SGLSTP &&
404         (pc == (greg_t)sys_sysenter || pc == (greg_t)brand_sys_sysenter ||
405          pc == (greg_t)tr_sys_sysenter ||
406          pc == (greg_t)tr_brand_sys_sysenter)) {
407     #else
408     if (tt == T_SGLSTP &&
409         (pc == (greg_t)sys_sysenter || pc == (greg_t)brand_sys_sysenter)) {
410     #endif
411         return (1);
412     }
413     return (0);
414 }
_____unchanged_portion_omitted_____
```