

```

*****
23565 Sun Nov  2 01:27:47 2014
new/usr/src/lib/libc/port/threads/sigaction.c
5281 incorrect realtime signal delivery
*****
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 2014 Ryan Zezeski. All rights reserved.
24 * Copyright 2010 Sun Microsystems, Inc. All rights reserved.
25 * Use is subject to license terms.
26 */

28 #include "lint.h"
29 #include <sys/feature_tests.h>
30 /*
31  * setcontext() really can return, if UC_CPU is not specified.
32  * Make the compiler shut up about it.
33  */
34 #if defined(__NORETURN)
35 #undef __NORETURN
36 #endif
37 #define __NORETURN
38 #include "thr_uberdata.h"
39 #include "asyncio.h"
40 #include <signal.h>
41 #include <siginfo.h>
42 #include <sys/system.h>

44 /* maskable signals */
45 const sigset_t maskset = {MASKSET0, MASKSET1, MASKSET2, MASKSET3};

47 /*
48  * Return true if the valid signal bits in both sets are the same.
49  */
50 int
51 sigequalset(const sigset_t *s1, const sigset_t *s2)
52 {
53     /*
54      * We only test valid signal bits, not rubbish following MAXSIG
55      * (for speed). Algorithm:
56      * if (s1 & fillset) == (s2 & fillset) then (s1 ^ s2) & fillset == 0
57      */
58     /* see lib/libc/inc/thr_uberdata.h for why this must be true */
59     #if (MAXSIG > (2 * 32) && MAXSIG <= (3 * 32))
60         return (!(s1->__sigbits[0] ^ s2->__sigbits[0]) |
61             (s1->__sigbits[1] ^ s2->__sigbits[1]) |

```

```

62         ((s1->__sigbits[2] ^ s2->__sigbits[2]) & FILLSET2));
63 #else
64 #error "fix me: MAXSIG out of bounds"
65 #endif
66 }

68 /*
69  * Common code for calling the user-specified signal handler.
70  */
71 void
72 call_user_handler(int sig, siginfo_t *sip, ucontext_t *ucp)
73 {
74     int i;
75     ulwp_t *self = curthread;
76     uberdata_t *udp = self->ul_uberdata;
77     struct sigaction uact;
78     volatile struct sigaction *sap;

80     /*
81      * If we are taking a signal while parked or about to be parked
82      * on __lwp_park() then remove ourselves from the sleep queue so
83      * that we can grab locks. The code in mutex_lock_queue() and
84      * cond_wait_common() will detect this and deal with it when
85      * __lwp_park() returns.
86      */
87     unsleep_self();
88     set_parking_flag(self, 0);

90     if (__td_event_report(self, TD_CATCHSIG, udp)) {
91         self->ul_td_evbuf.eventnum = TD_CATCHSIG;
92         self->ul_td_evbuf.eventdata = (void *) (intptr_t) sig;
93         tdb_event(TD_CATCHSIG, udp);
94     }

96     /*
97      * Get a self-consistent set of flags, handler, and mask
98      * while holding the sig's sig_lock for the least possible time.
99      * We must acquire the sig's sig_lock because some thread running
100     * in sigaction() might be establishing a new signal handler.
101     * The code in sigaction() acquires the writer lock; here
102     * we acquire the readers lock to enhance concurrency in the
103     * face of heavy signal traffic, such as generated by java.
104     */
105     /* Locking exceptions:
106     * No locking for a child of vfork().
107     * If the signal is SIGPROF with an si_code of PROF_SIG,
108     * then we assume that this signal was generated by
109     * setitimer(ITIMER_REALPROF) set up by the dbx collector.
110     * If the signal is SIGEMT with an si_code of EMT_CPCOVF,
111     * then we assume that the signal was generated by
112     * a hardware performance counter overflow.
113     * In these cases, assume that we need no locking. It is the
114     * monitoring program's responsibility to ensure correctness.
115     */
116     sap = &udp->sigaction[sig].sig_uaction;
117     if (self->ul_vfork ||
118         (sip != NULL &&
119          ((sig == SIGPROF && sip->si_code == PROF_SIG) ||
120           (sig == SIGEMT && sip->si_code == EMT_CPCOVF)))) {
121         /* we wish this assignment could be atomic */
122         (void) memcpy(&uact, (void *) sap, sizeof (uact));
123     } else {
124         rwlock_t *rwlp = &udp->sigaction[sig].sig_lock;
125         lrw_rdlock(rwlp);
126         (void) memcpy(&uact, (void *) sap, sizeof (uact));
127         if ((sig == SIGCANCEL || sig == SIGAIOCANCEL) &&

```

```

128         (sap->sa_flags & SA_RESETHAND))
129         sap->sa_sigaction = SIG_DFL;
130         lrw_unlock(rwlp);
131     }
132
133     /*
134     * Set the proper signal mask and call the user's signal handler.
135     * (We override the user-requested signal mask with maskset
136     * so we currently have all blockable signals blocked.)
137     *
138     * We would like to ASSERT() that the signal is not a member of the
139     * signal mask at the previous level (ucp->uc_sigmask) or the specified
140     * signal mask for sigsuspend() or pollsys() (self->ul_tmpmask) but
141     * /proc can override this via PCSSIG, so we don't bother.
142     *
143     * We would also like to ASSERT() that the signal mask at the previous
144     * level equals self->ul_sigmask (maskset for sigsuspend() / pollsys()),
145     * but /proc can change the thread's signal mask via PCSHOLD, so we
146     * don't bother with that either.
147     */
148     ASSERT(ucp->uc_flags & UC_SIGMASK);
149     if (self->ul_sigsuspend) {
150         ucp->uc_sigmask = self->ul_sigmask;
151         self->ul_sigsuspend = 0;
152         /* the sigsuspend() or pollsys() signal mask */
153         sigorset(&uact.sa_mask, &self->ul_tmpmask);
154     } else {
155         /* the signal mask at the previous level */
156         sigorset(&uact.sa_mask, &ucp->uc_sigmask);
157     }
158     if (!(uact.sa_flags & SA_NODEFER)) /* add current signal */
159         (void) sigaddset(&uact.sa_mask, sig);
160
161     /*
162     * Enforce the proper order for realtime signals. Lower signals
163     * have higher priority and multiple instances of the same signal
164     * must arrive in FIFO order (NODEFER does not apply).
165     *
166     * See section 2.4.2 of POSIX.
167     */
168     if ((sig >= SIGRTMIN) && (sig <= SIGRTMAX)) {
169         for (i = sig; i <= SIGRTMAX; i++) {
170             (void) sigaddset(&uact.sa_mask, i);
171         }
172     }
173
174     self->ul_sigmask = uact.sa_mask;
175     self->ul_siglink = ucp;
176     (void) __lwp_sigmask(SIG_SETMASK, &uact.sa_mask);
177
178     /*
179     * If this thread has been sent SIGCANCEL from the kernel
180     * or from pthread_cancel(), it is being asked to exit.
181     * The kernel may send SIGCANCEL without a siginfo struct.
182     * If the SIGCANCEL is process-directed (from kill() or
183     * sigqueue()), treat it as an ordinary signal.
184     */
185     if (sig == SIGCANCEL) {
186         if (sip == NULL || SI_FROMKERNEL(sip) ||
187             sip->si_code == SI_LWP) {
188             do_sigcancel();
189             goto out;
190         }
191         /* SIGCANCEL is ignored by default */
192         if (uact.sa_sigaction == SIG_DFL ||
193             uact.sa_sigaction == SIG_IGN)

```

```

194         goto out;
195     }
196
197     /*
198     * If this thread has been sent SIGAIOCANCEL (SIGLWP) and
199     * we are an aio worker thread, cancel the aio request.
200     */
201     if (sig == SIGAIOCANCEL) {
202         aio_worker_t *aiowp = pthread_getspecific(_aio_key);
203
204         if (sip != NULL && sip->si_code == SI_LWP && aiowp != NULL)
205             siglongjmp(aiowp->work_jmp_buf, 1);
206         /* SIGLWP is ignored by default */
207         if (uact.sa_sigaction == SIG_DFL ||
208             uact.sa_sigaction == SIG_IGN)
209             goto out;
210     }
211
212     if (!(uact.sa_flags & SA_SIGINFO))
213         sip = NULL;
214     __sighndlr(sig, sip, ucp, uact.sa_sigaction);
215
216 #if defined(sparc) || defined(__sparc)
217     /*
218     * If this is a floating point exception and the queue
219     * is non-empty, pop the top entry from the queue. This
220     * is to maintain expected behavior.
221     */
222     if (sig == SIGFPE && ucp->uc_mcontext.fpregs.fpu_qcnt) {
223         fpregset_t *fp = &ucp->uc_mcontext.fpregs;
224
225         if (--fp->fpu_qcnt > 0) {
226             unsigned char i;
227             struct fq *fq;
228
229             fq = fp->fpu_q;
230             for (i = 0; i < fp->fpu_qcnt; i++)
231                 fq[i] = fq[i+1];
232         }
233     }
234 #endif /* sparc */
235
236 out:
237     (void) setcontext(ucp);
238     thr_panic("call_user_handler(): setcontext() returned");
239 }

```

unchanged portion omitted

```

*****
8336 Sun Nov  2 01:27:47 2014
new/usr/src/man/man2/sigaction.2
5281 incorrect realtime signal delivery
*****
1 \" te
2 .\" Copyright (c) 2014, Ryan Zezeski. All Rights Reserved.
3 .\" Copyright (c) 2007, Sun Microsystems, Inc. All Rights Reserved.
4 .\" Copyright 1989 AT&T
5 .\" The contents of this file are subject to the terms of the Common Development
6 .\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE or http:
7 .\" When distributing Covered Code, include this CDDL HEADER in each file and in
8 .TH SIGACTION 2 "Oct 22, 2014"
9 .TH SIGACTION 2 "Mar 23, 2005"
10 .SH NAME
11 sigaction \- detailed signal management
12 .SH SYNOPSIS
13 .LP
14 #include <signal.h>
15
16 \fBint\fR \fBsigaction\fR(\fBint\fR \fBisig\fR, \fBconst struct sigaction *restrictri
17 \fBstruct sigaction *restrict\fR \fBioact\fR);
18 .fi
19
20 .SH DESCRIPTION
21 .sp
22 .LP
23 The \fBsigaction()\fR function allows the calling process to examine or specify
24 the action to be taken on delivery of a specific signal. See
25 \fBsignal.h\fR(3HEAD) for an explanation of general signal concepts.
26 .sp
27 .LP
28 The \fBisig\fR argument specifies the signal and can be assigned any of the
29 signals specified in \fBsignal.h\fR(3HEAD) except \fBBSIGKILL\fR and
30 \fBBSIGSTOP\fR.
31 .sp
32 .LP
33 If the argument \fBioact\fR is not \fBINULL\fR, it points to a structure
34 specifying the new action to be taken when delivering \fBisig\fR. If the
35 argument \fBioact\fR is not \fBINULL\fR, it points to a structure where the
36 action previously associated with \fBisig\fR is to be stored on return from
37 \fBsigaction()\fR.
38 .sp
39 .LP
40 The \fBsigaction\fR structure includes the following members:
41 .sp
42 .in +2
43 void (*sa_handler)();
44 void (*sa_sigaction)(int, siginfo_t *, void *);
45 sigset_t sa_mask;
46 int sa_flags;
47 .fi
48 .in -2
49
50 .sp
51 .LP
52 The storage occupied by \fBsa_handler\fR and \fBsa_sigaction\fR may overlap,
53 and a standard-conforming application (see \fBstandards\fR(5)) must not use
54 both simultaneously.
55 .sp
56 .LP
57 The \fBsa_handler\fR member identifies the action to be associated with the
58 specified signal, if the \fBBSA_SIGINFO\fR flag (see below) is cleared in the
59 \fBsa_flags\fR field of the sigaction structure. It may take any of the values

```

```

60 specified in \fBsignal.h\fR(3HEAD) or that of a user specified signal handler.
61 If the \fBBSA_SIGINFO\fR flag is set in the \fBsa_flags\fR field, the
62 \fBsa_sigaction\fR field specifies a signal-catching function.
63 .sp
64 .LP
65 The \fBsa_mask\fR member specifies a set of signals to be blocked while the
66 signal handler is active. On entry to the signal handler, that set of signals
67 is added to the set of signals already being blocked when the signal is
68 delivered. In addition, the signal that caused the handler to be executed will
69 also be blocked, unless the \fBBSA_NODEFER\fR flag has been specified.
70 \fBBSIGSTOP\fR and \fBBSIGKILL\fR cannot be blocked (the system silently
71 enforces this restriction).
72 .sp
73 .LP
74 The \fBsa_flags\fR member specifies a set of flags used to modify the delivery
75 of the signal. It is formed by a logical \fBOR\fR of any of the following
76 values:
77 .sp
78 .ne 2
79 .na
80 \fBBSA_ONSTACK\fR
81 .ad
82 .RS 16n
83 If set and the signal is caught, and if the thread that is chosen to processes
84 a delivered signal has an alternate signal stack declared with
85 \fBsigaltstack\fR(2), then it will process the signal on that stack. Otherwise,
86 the signal is delivered on the thread's normal stack.
87 .RE
88
89 .sp
90 .ne 2
91 .na
92 \fBBSA_RESETHAND\fR
93 .ad
94 .RS 16n
95 If set and the signal is caught, the disposition of the signal is reset to
96 \fBBSIG_DFL\fR and the signal will not be blocked on entry to the signal handler
97 (\fBBSIGILL\fR, \fBBSIGTRAP\fR, and \fBBSIGPWR\fR cannot be automatically reset
98 when delivered; the system silently enforces this restriction).
99 .RE
100
101 .sp
102 .ne 2
103 .na
104 \fBBSA_NODEFER\fR
105 .ad
106 .RS 16n
107 If set and the signal is caught, the signal will not be automatically
108 blocked by the kernel while it is being caught. Unless the signal is
109 a realtime signal. Multiple instances of the same realtime signal
110 must be delivered in FIFO order and thus are always deferred.
111 If set and the signal is caught, the signal will not be automatically blocked
112 by the kernel while it is being caught.
113 .RE
114
115 .sp
116 .ne 2
117 .na
118 \fBBSA_RESTART\fR
119 .ad
120 .RS 16n
121 If set and the signal is caught, functions that are interrupted by the
122 execution of this signal's handler are transparently restarted by the system,
123 namely \fBfcntl\fR(2), \fBioctl\fR(2), \fBwait\fR(3C), \fBwaitid\fR(2), and the
124 following functions on slow devices like terminals: \fBgetmsg()\fR and
125 \fBgetpmsg()\fR (see \fBgetmsg\fR(2)); \fBputmsg()\fR and \fBputpmsg()\fR (see

```

```

124 \fBputmsg\fR(2)); \fBpread()\fR, \fBread()\fR, and \fBreadv()\fR (see
125 \fBread\fR(2)); \fBpwrite()\fR, \fBwrite()\fR, and \fBwritev()\fR (see
126 \fBwrite\fR(2)); \fBrecv()\fR, \fBrecvfrom()\fR, and \fBrecvmsg()\fR (see
127 \fBrecv\fR(3SOCKET)); and \fBsend()\fR, \fBsendto()\fR, and \fBsendmsg()\fR
128 (see \fBsend\fR(3SOCKET)). Otherwise, the function returns an \fBEINTR\fR
129 error.
130 .RE

132 .sp
133 .ne 2
134 .na
135 \fB\fBSA_SIGINFO\fR\fR
136 .ad
137 .RS 16n
138 If cleared and the signal is caught, \fIisig\fR is passed as the only argument
139 to the signal-catching function. If set and the signal is caught, two
140 additional arguments are passed to the signal-catching function. If the second
141 argument is not equal to \fINULL\fR, it points to a \fBsiginfo_t\fR structure
142 containing the reason why the signal was generated (see
143 \fBsiginfo.h\fR(3HEAD)); the third argument points to a \fBucontext_t\fR
144 structure containing the receiving process's context when the signal was
145 delivered (see \fBucontext.h\fR(3HEAD)).
146 .RE

148 .sp
149 .ne 2
150 .na
151 \fB\fBSA_NOCLDWAIT\fR\fR
152 .ad
153 .RS 16n
154 If set and \fIisig\fR equals \fBSIGCHLD\fR, the system will not create zombie
155 processes when children of the calling process exit. If the calling process
156 subsequently issues a \fBwait\fR(3C), it blocks until all of the calling
157 process's child processes terminate, and then returns \fB\{mil\fR with
158 \fBerrno\fR set to \fBECHILD\fR.
159 .RE

161 .sp
162 .ne 2
163 .na
164 \fB\fBSA_NOCLDSTOP\fR\fR
165 .ad
166 .RS 16n
167 If set and \fIisig\fR equals \fBSIGCHLD\fR, \fBSIGCHLD\fR will not be sent to
168 the calling process when its child processes stop or continue.
169 .RE

171 .SH RETURN VALUES
170 .sp
172 .LP
173 Upon successful completion, \fB0\fR is returned. Otherwise, \fB\{mil\fR is
174 returned, \fBerrno\fR is set to indicate the error, and no new signal handler
175 is installed.
176 .SH ERRORS
176 .sp
177 .LP
178 The \fBsigaction()\fR function will fail if:
179 .sp
180 .ne 2
181 .na
182 \fB\fBEINVAL\fR\fR
183 .ad
184 .RS 10n
185 The value of the \fIisig\fR argument is not a valid signal number or is equal to
186 \fBSIGKILL\fR or \fBSIGSTOP\fR. In addition, if in a multithreaded process, it
187 is equal to \fBSIGWAITING\fR, \fBSIGCANCEL\fR, or \fBSIGLWP\fR.

```

```

188 .RE

190 .SH ATTRIBUTES
191 .sp
192 See \fBattributes\fR(5) for descriptions of the following attributes:
193 .sp

195 .sp
196 .TS
197 box;
198 c | c
199 l | l .
200 ATTRIBUTE TYPE ATTRIBUTE VALUE
201 _
202 Interface Stability Committed
203 _
204 MT-Level Async-Signal-Safe
205 _
206 Standard See \fBstandards\fR(5).
207 .TE

209 .SH SEE ALSO
210 .sp
211 .LP
212 \fBkill\fR(1), \fBintro\fR(2), \fBexit\fR(2), \fBfcntl\fR(2), \fBgetmsg\fR(2),
213 \fBioctl\fR(2), \fBkill\fR(2), \fBpause\fR(2), \fBputmsg\fR(2), \fBread\fR(2),
214 \fBsigaltstack\fR(2), \fBsigprocmask\fR(2), \fBsigsend\fR(2),
215 \fBsigsuspend\fR(2), \fBwaitid\fR(2), \fBwrite\fR(2), \fBrecv\fR(3SOCKET),
216 \fBsend\fR(3SOCKET), \fBsiginfo.h\fR(3HEAD), \fBsignal\fR(3C),
217 \fBsignal.h\fR(3HEAD), \fBsigsetops\fR(3C), \fBucontext.h\fR(3HEAD),
218 \fBwait\fR(3C), \fBattributes\fR(5), \fBstandards\fR(5)
219 .SH NOTES
220 The handler routine can be declared:
221 .sp
222 .in +2
223 .nf
224 void handler (int \fIisig\fR, siginfo_t *\fIisip\fR, ucontext_t *\fIiucp\fR);
225 .fi
226 .in -2

228 .sp
229 .LP
230 The \fIisig\fR argument is the signal number. The \fIisip\fR argument is a
231 pointer (to space on the stack) to a \fBsiginfo_t\fR structure, which provides
232 additional detail about the delivery of the signal. The \fIiucp\fR argument is a
233 pointer (again to space on the stack) to a \fBucontext_t\fR structure (defined
234 in <\fBsys/ucontext.h\fR>) which contains the context from before the signal.
235 It is not recommended that \fIiucp\fR be used by the handler to restore the
236 context from before the signal delivery.

```

new/usr/src/pkg/manifests/system-test-ostest.mf

1

\*\*\*\*\*

1606 Sun Nov 2 01:27:48 2014

new/usr/src/pkg/manifests/system-test-ostest.mf

5281 incorrect realtime signal delivery

\*\*\*\*\*

1 #  
2 # This file and its contents are supplied under the terms of the  
3 # Common Development and Distribution License ("CDDL"), version 1.0.  
4 # You may only use this file in accordance with the terms of version  
5 # 1.0 of the CDDL.  
6 #  
7 # A full copy of the text of the CDDL should have accompanied this  
8 # source. A copy of the CDDL is also available via the Internet at  
9 # <http://www.illumos.org/license/CDDL>.  
10 #

12 #  
13 # Copyright (c) 2012 by Delphix. All rights reserved.  
14 # Copyright 2014, OmniTI Computer Consulting, Inc. All rights reserved.  
15 # **Copyright 2014, Ryan Zezeski. All rights reserved.**  
16 #

18 set name=pkg.fmri value=pkg:/system/test/ostest@\$(PKGVERS)  
19 set name=pkg.description value="Miscellaneous OS Unit Tests"  
20 set name=pkg.summary value="OS Unit Test Suite"  
21 set name=info.classification \  
22 value=org.opensolaris.category.2008:Development/System  
23 set name=variant.arch value=\$(ARCH)  
24 dir path=opt/os-tests  
25 dir path=opt/os-tests/bin  
26 dir path=opt/os-tests/runfiles  
27 dir path=opt/os-tests/tests  
28 **dir path=opt/os-tests/tests/rt\_delivery**  
29 dir path=opt/os-tests/tests/sigqueue  
30 file path=opt/os-tests/README mode=0444  
31 file path=opt/os-tests/bin/ostest mode=0555  
32 file path=opt/os-tests/runfiles/delphix.run mode=0444  
33 file path=opt/os-tests/runfiles/omnios.run mode=0444  
34 file path=opt/os-tests/runfiles/openindiana.run mode=0444  
35 file path=opt/os-tests/tests/poll\_test mode=0555  
36 **file path=opt/os-tests/tests/rt\_delivery/rt\_delivery mode=0555**  
37 file path=opt/os-tests/tests/sigqueue/sigqueue\_queue\_size mode=0555  
38 license cr\_Sun license=cr\_Sun  
39 license lic\_CDDL license=lic\_CDDL  
40 depend fmri=system/test/testrunner type=require

new/usr/src/test/os-tests/runfiles/delphix.run

1

\*\*\*\*\*

789 Sun Nov 2 01:27:48 2014

new/usr/src/test/os-tests/runfiles/delphix.run

5281 incorrect realtime signal delivery

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet at
9 # http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright (c) 2014 by Ryan Zezeski. All rights reserved.
14 # Copyright (c) 2012 by Delphix. All rights reserved.
15 #
```

```
17 [DEFAULT]
18 pre =
19 verbose = False
20 quiet = False
21 timeout = 60
22 post =
23 outputdir = /var/tmp/test_results
```

```
25 [/opt/os-tests/tests/poll_test]
26 user = root
```

```
28 [/opt/os-tests/tests/sigqueue]
29 tests = ['sigqueue_queue_size']
```

```
31 [/opt/os-tests/tests/rt_delivery]
32 tests = ['rt_delivery']
```

new/usr/src/test/os-tests/runfiles/omnios.run

1

\*\*\*\*\*

789 Sun Nov 2 01:27:48 2014

new/usr/src/test/os-tests/runfiles/omnios.run

5281 incorrect realtime signal delivery

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet at
9 # http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright (c) 2014 by Ryan Zezeski. All rights reserved.
14 # Copyright (c) 2012 by Delphix. All rights reserved.
15 #
```

```
17 [DEFAULT]
18 pre =
19 verbose = False
20 quiet = False
21 timeout = 60
22 post =
23 outputdir = /var/tmp/test_results
```

```
25 [/opt/os-tests/tests/poll_test]
26 user = root
```

```
28 [/opt/os-tests/tests/sigqueue]
29 tests = ['sigqueue_queue_size']
```

```
31 [/opt/os-tests/tests/rt_delivery]
32 tests = ['rt_delivery']
```

new/usr/src/test/os-tests/runfiles/openindiana.run

1

\*\*\*\*\*

789 Sun Nov 2 01:27:48 2014

new/usr/src/test/os-tests/runfiles/openindiana.run

5281 incorrect realtime signal delivery

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet at
9 # http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright (c) 2014 by Ryan Zezeski. All rights reserved.
14 # Copyright (c) 2012 by Delphix. All rights reserved.
15 #
```

```
17 [DEFAULT]
18 pre =
19 verbose = False
20 quiet = False
21 timeout = 60
22 post =
23 outputdir = /var/tmp/test_results
```

```
25 [/opt/os-tests/tests/poll_test]
26 user = root
```

```
28 [/opt/os-tests/tests/sigqueue]
29 tests = ['sigqueue_queue_size']
```

```
31 [/opt/os-tests/tests/rt_delivery]
32 tests = ['rt_delivery']
```



new/usr/src/test/os-tests/tests/Makefile

1

\*\*\*\*\*

591 Sun Nov 2 01:27:48 2014

new/usr/src/test/os-tests/tests/Makefile

5281 incorrect realtime signal delivery

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet at
9 # http://www.illumos.org/license/CDDL.
10 #
```

```
12 #
13 # Copyright (c) 2014 by Ryan Zezeski. All rights reserved.
14 # Copyright (c) 2012 by Delphix. All rights reserved.
15 #
```

```
17 SUBDIRS = poll sigqueue rt_delivery
16 SUBDIRS = poll sigqueue
```

```
19 include $(SRC)/test/Makefile.com
```

new/usr/src/test/os-tests/tests/rt\_delivery/Makefile

1

\*\*\*\*\*

1043 Sun Nov 2 01:27:48 2014

new/usr/src/test/os-tests/tests/rt\_delivery/Makefile

5281 incorrect realtime signal delivery

\*\*\*\*\*

```
1 #
2 # This file and its contents are supplied under the terms of the
3 # Common Development and Distribution License ("CDDL"), version 1.0.
4 # You may only use this file in accordance with the terms of version
5 # 1.0 of the CDDL.
6 #
7 # A full copy of the text of the CDDL should have accompanied this
8 # source. A copy of the CDDL is also available via the Internet at
9 # http://www.illumos.org/license/CDDL.
10 #
11 #
12 #
13 # Copyright (c) 2014 Ryan Zezeski. All rights reserved.
14 #
15 #
16 include $(SRC)/cmd/Makefile.cmd
17 include $(SRC)/test/Makefile.com
18 #
19 PROG = rt_delivery
20 OBJS = $(PROG:%=%.o)
21 SRCS = $(OBJS:%.o=%.c)
22 #
23 C99MODE = -xc99=%all
24 #
25 ROOTOPTPKG = $(ROOT)/opt/os-tests
26 TESTDIR = $(ROOTOPTPKG)/tests/rt_delivery
27 #
28 CMDS = $(PROG:%=$(TESTDIR)/%)
29 $(CMDS) := FILEMODE = 0555
30 #
31 all: $(PROG)
32 #
33 $(PROG): $(OBJS)
34     $(LINK.c) $(OBJS) -o $@ $(LDLIBS)
35     $(POST_PROCESS)
36 #
37 %.o: ../%.c
38     $(COMPILE.c) $<
39 #
40 install: all $(CMDS)
41 #
42 lint: lint_SRCS
43 #
44 clobber: clean
45     -$(RM) $(PROG)
46 #
47 clean:
48     -$(RM) $(OBJS)
49 #
50 $(CMDS): $(TESTDIR) $(PROG)
51 #
52 $(TESTDIR):
53     $(INS.dir)
54 #
55 $(TESTDIR)/%: %
56     $(INS.file)
```

new/usr/src/test/os-tests/tests/rt\_delivery/rt\_delivery.c

1

```
*****
6015 Sun Nov  2 01:27:48 2014
new/usr/src/test/os-tests/tests/rt_delivery/rt_delivery.c
5281 incorrect realtime signal delivery
*****
1 /*
2  * This file and its contents are supplied under the terms of the
3  * Common Development and Distribution License ("CDDL"), version 1.0.
4  * You may only use this file in accordance with the terms of version
5  * 1.0 of the CDDL.
6  *
7  * A full copy of the text of the CDDL should have accompanied this
8  * source.  A copy of the CDDL is also available via the Internet at
9  * http://www.illumos.org/license/CDDL.
10 *
11 * Copyright 2014 Ryan Zezeski.  All rights reserved.
12 */
13 #include <signal.h>
14 #include <stdarg.h>
15 #include <stdio.h>
16 #include <stdlib.h>
17 #include <sys/types.h>
18 #include <sys/wait.h>
19 #include <unistd.h>
20
21 /*
22  * The 'seq' struct and 'sh' handler are used to record the order in which
23  * the signals are delivered to the parent process.  When a signal is
24  * delivered the handler records it in seq and then increments the index.
25  */
26 struct seq {
27     int sig, num;
28 };
29
30 static struct seq *seqp;
31
32 /* ARGSUSED */
33 void
34 sh(int sig, siginfo_t *info, void *unused)
35 {
36     static int i;
37     seqp[i].num = info->si_value.sival_int;
38     seqp[i++].sig = sig;
39 }
40
41 static void
42 test_start(const char *test_name, const char *format, ...)
43 {
44     va_list args;
45
46     (void) printf("TEST STARTING %s: ", test_name);
47
48     va_start(args, format);
49     (void) vprintf(format, args);
50     va_end(args);
51     (void) fflush(stdout);
52 }
53
54 static void
55 test_failed(const char *test_name, const char *format, ...)
56 {
57     va_list args;
58
59     (void) printf("TEST FAILED %s: ", test_name);
60
61     va_start(args, format);
```

new/usr/src/test/os-tests/tests/rt\_delivery/rt\_delivery.c

2

```
62     (void) vprintf(format, args);
63     va_end(args);
64
65     (void) exit(-1);
66 }
67
68 static void
69 test_passed(const char *test_name)
70 {
71     (void) printf("TEST PASS: %s\n", test_name);
72     (void) fflush(stdout);
73 }
74
75 /*
76  * Verify the delivery order of realtime (RT) signals.  RT signals were
77  * introduced in POSIX.1b and differ from regular signals in several key
78  * ways:
79  *
80  * 1. Multiple instances of the same signal are queued and must be
81  * delivered in FIFO order.
82  *
83  * 2. If multiple different RT signals are pending then the lower numbered
84  * RT signals are delivered first.  I.e., RT signals have a notion of
85  * priority.
86  *
87  * See sections 2.4.2 (Realtime Signal Generation and Delivery) and B.2.4
88  * of the POSIX.1-2008 standard for more information.
89  */
90 static int
91 rt_delivery(const char *test_name, int no_defer)
92 {
93     int i, j, sig, sival, sigend, numsig, rc, fail = 0;
94     pid_t p = getpid(), chld;
95     struct sigaction sa;
96     sigset_t ss, ssold;
97     union sigval sv;
98
99     /*
100    * Range of signals to send.  Each signal is sent 2 times to verify FIFO
101    * order.
102    */
103     sigend = SIGRTMIN + (sysconf(_SC_SIGQUEUE_MAX) / 2);
104     sigend = sigend > SIGRTMAX ? SIGRTMAX : sigend;
105     numsig = (sigend - SIGRTMIN) * 2;
106
107     seqp = malloc(sizeof (struct seq) * numsig);
108
109     /*
110    * Step 1:
111    *
112    * Install the signal handler (sh) on this process (the parent).
113    */
114     (void) sigemptyset(&ss);
115     for (i = SIGRTMIN; i < sigend; i++) {
116         sa.sa_sigaction = sh;
117         sa.sa_flags = SA_SIGINFO;
118         sa.sa_flags |= (no_defer ? SA_NODEFER : 0);
119         (void) sigemptyset(&sa.sa_mask);
120         sigaction(i, &sa, NULL);
121         (void) sigaddset(&ss, i);
122     }
123
124     /*
125    * Step 2:
126    *
127    * Block all signals that are about to be sent to this process.
```

```

128  * This is needed to make sure all signals are pending at once.
129  * This is related to step 4.
130  */
131  (void) sigprocmask(SIG_BLOCK, &ss, &ssold);

133  /*
134  * Step 3:
135  *
136  * Fork a child process to send each RT signal 2 times to the
137  * parent process. Remember, all signals being sent are RT
138  * signals and thus are subject to explicit ordering guarantees
139  * spelled out in POSIX.1b. Lower numbered RT signals are
140  * _always_ delivered before higher RT signals.
141  */
142  if ((chld = fork()) == 0) {
143      for (i = SIGRTMIN; i < sigend; i++) {
144          for (j = 0; j < 2; j++) {
145              sv.sival_int = j;
146              if (sigqueue(p, i, sv) == -1) {
147                  perror("problem generating signal");
148                  exit(1);
149              }
150          }
151      }
152      exit(0);
153  }

155  /*
156  * Step 4:
157  *
158  * Wait for forked process to finish sending all signals. This is
159  * to make sure that all signals are pending before allowing the
160  * kernel to deliver them to the handler. If you don't both wait
161  * and block then the handler will process the signals as they
162  * come in and thus we will not be testing the ordering
163  * guarantees. Related to step 2.
164  */
165  if (wait(&rc) != chld) {
166      test_failed(test_name, "failed to wait for child\n");
167  }

169  if (rc != 0) {
170      test_failed(test_name, "child exited non-zero\n");
171  }
172  (void) sigprocmask(SIG_SETMASK, &ssold, NULL);

174  /*
175  * Step 5:
176  *
177  * Verify that multiple pending RT signals are delivered in
178  * correct priority (lower signals first), and that multiple
179  * instances of the same RT signal are delivered in FIFO order.
180  *
181  * E.g. if SIGRTMIN is 42 then the handler should record this
182  * ordering:
183  *
184  * {42, 0}, {42, 1}, {42, 2}, {43, 0}, {43, 1}, {43, 2}, ...
185  */
186  sig = SIGRTMIN;
187  sival = 0;
188  for (i = 0; i < numsig; i++) {
189      if ((seqp[i].sig != sig) || (seqp[i].num != sival)) {
190          fail = 1;
191      }
192      sig = ((i + 1) % 2) == 0 ? sig + 1 : sig;
193      sival = ((i + 1) % 2) == 0 ? 0 : sival + 1;

```

```

194  }
195  }
196  if (fail) {
197      sig = SIGRTMIN;
198      sival = 0;
199      (void) fprintf(stderr, "Excpeted Actual\n");
200      for (i = 0; i < numsig; i++) {
201          (void) fprintf(stderr, "{%2d,%d}  {%2d,%d}\n",
202                          sig, sival, seqp[i].sig, seqp[i].num);
203      }
204      sig = ((i + 1) % 2) == 0 ? sig + 1 : sig;
205      sival = ((i + 1) % 2) == 0 ? 0 : sival + 1;
206  }
207  test_failed(test_name, "incorrect order\n");
208  }

210  free(seqp);
211  return (0);
212  }

214  static void
215  rt_delivery_defer_test(void)
216  {
217      const char *test_name = __func__;
218      test_start(test_name, "verify delivery order of pending RT signals\n");
219      (void) rt_delivery(test_name, 0);
220      test_passed(test_name);
221  }

223  static void
224  rt_delivery_nodfer_test(void)
225  {
226      const char *test_name = __func__;
227      test_start(test_name,
228                "verify delivery order of pending RT signals w/ SA_NODEFER\n");
229      (void) rt_delivery(test_name, 1);
230      test_passed(test_name);
231  }

233  static void
234  run_tests(void)
235  {
236      rt_delivery_defer_test();
237      rt_delivery_nodfer_test();
238  }

240  /* ARGSUSED */
241  int
242  main(int argc, char *argv[])
243  {
244      run_tests();
245      return (EXIT_SUCCESS);
246  }

```