

new/usr/src/cmd/utmpd/svc-utmpd

1

```
*****
1307 Mon Jan 12 23:37:07 2015
new/usr/src/cmd/utmpd/svc-utmpd
5375 utmpd(1M) core dumps when WTMPX_UPDATE_FREQ is zero
*****
1 #!/sbin/sh
2 #
3 # CDDL HEADER START
4 #
5 # The contents of this file are subject to the terms of the
6 # Common Development and Distribution License (the "License").
7 # You may not use this file except in compliance with the License.
8 #
9 # You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
10 # or http://www.opensolaris.org/os/licensing.
11 # See the License for the specific language governing permissions
12 # and limitations under the License.
13 #
14 # When distributing Covered Code, include this CDDL HEADER in each
15 # file and include the License file at usr/src/OPENSOLARIS.LICENSE.
16 # If applicable, add the following below this CDDL HEADER, with the
17 # fields enclosed by brackets "[]" replaced with your own identifying
18 # information: Portions Copyright [yyyy] [name of copyright owner]
19 #
20 # CDDL HEADER END
21 #
22 #
23 # Copyright 2007 Sun Microsystems, Inc. All rights reserved.
24 # Use is subject to license terms.
25 #
26 # ident "%Z%M% %I% %E% SMI"

27 . /lib/svc/share/smf_include.sh

29 [ ! -x /usr/lib/utmpd ] && exit $SMF_EXIT_ERR_CONFIG

31 # If a utmppipe exists, check for a utmpd process and exit
32 # if the daemon is already running.

34 if [ -p /var/run/utmppipe ]; then
35     if /usr/bin/pgrep -x -u 0 -z `smf_zonename` utmpd >/dev/null 2>&1; then
36         echo "$0: utmpd is already running"
37         exit 1
38     fi
39 fi

41 /usr/bin/rm -f /var/run/utmppipe
42 /usr/lib/utmpd &
```

```

*****
26383 Mon Jan 12 23:37:07 2015
new/usr/src/cmd/utmpd/utmpd.c
5375 utmpd(1M) core dumps when WTMPX_UPDATE_FREQ is zero
*****
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 2014, 2015 Shruti V Sampat <shrutisampat@gmail.com>
23 * Copyright 2014 Shruti V Sampat <shrutisampat@gmail.com>
24 */
25 /*
26 * Copyright 2007 Sun Microsystems, Inc. All rights reserved.
27 * Use is subject to license terms.
28 */
29
30 /*      Copyright (c) 1984, 1986, 1987, 1988, 1989 AT&T */
31 /*      All Rights Reserved      */
32
33 /*
34 * Portions of such source code were derived from Berkeley 4.3 BSD
35 * under license from the Regents of the University of California.
36 */
37
38 /*
39 * utmpd          - utmp daemon
40 *
41 * This program receives requests from pututxline(3)
42 * via a named pipe to watch the process to make sure it cleans up
43 * its utmpx entry on termination.
44 * The program keeps a list of procs
45 * and uses poll() on their /proc files to detect termination.
46 * Also the program periodically scans the /etc/utmpx file for
47 * processes that aren't in the table so they can be watched.
48 *
49 * If utmpd doesn't hear back over the pipe from pututline(3) that
50 * the process has removed its entry it cleans the entry when the
51 * process terminates.
52 * The AT&T Copyright above is there since we borrowed the pipe
53 * mechanism from init(1m).
54 */
55
56 #include <sys/types.h>
57 #include <signal.h>
58 #include <stdio.h>
59 #include <stdio_ext.h>

```

```

61 #include <unistd.h>
62 #include <utmpx.h>
63 #include <errno.h>
64 #include <termio.h>
65 #include <sys/termios.h>
66 #include <sys/tty.h>
67 #include <ctype.h>
68 #include <sys/stat.h>
69 #include <sys/statvfs.h>
70 #include <fcntl.h>
71 #include <time.h>
72 #include <sys/stropts.h>
73 #include <wait.h>
74 #include <syslog.h>
75 #include <stdlib.h>
76 #include <string.h>
77 #include <poll.h>
78 #include <deflt.h>
79 #include <proefs.h>
80 #include <sys/resource.h>
81 #include <limits.h>
82
83 #define dprintf(x)      if (Debug) (void) printf x
84
85 /*
86  * Memory allocation keyed off MAX_FDS
87  */
88 #define MAX_FDS          4064      /* Maximum # file descriptors */
89 #define EXTRA_MARGIN    32        /* Allocate this many more FDS over Max_Fds */
90 /*
91  * MAX_POLLNV & RESETS - paranoia to cover an error case that might not exist
92  */
93 #define MAX_POLL_ERRS    1024     /* Count of bad errors */
94 #define MAX_RESETS       1024     /* Maximum times to reload tables */
95 #define POLL_TIMEOUT     300      /* Default Timeout for poll() in seconds */
96 #define CLEANIT          1        /* Used by rem_pid() */
97 #define DONT_CLEAN       0        /* Used by rem_pid() */
98 #define UTMP_DEFAULT     "/etc/default/utmpd"
99 #define WARN_TIME        3600     /* seconds between utmp checks */
100 #define WTMPX_UFREQ      60       /* seconds between updating WTMPX's atime */
101
102
103 /*
104 * The pidrec structure describes the data shipped down the pipe to
105 * us from the pututxline() library in
106 * lib/libc/port/gen/getutx.c
107 */
108
109 /*
110  * pd_type's
111  */
112 #define ADDPID 1
113 #define REMPID 2
114
115 struct pidrec {
116     int      pd_type;          /* Command type */
117     pid_t    pd_pid;          /* pid to add or remove */
118 };
119
120 unchanged_portion_omitted
121
122 static struct pidentry *pidtable = NULL;
123
124 static pollfd_t *fdtable = NULL;
125
126 static int      pidcnt = 0;    /* Number of procs being watched */
127 static char     *prog_name;    /* To save the invocation name away */

```

```

140 static char *UTMPPIPE_DIR = "/var/run";
141 static char *UTMPPIPE = "/var/run/utmppipe";
142 static int Pfd = -1; /* File descriptor of named pipe */
143 static int Poll_timeout = POLL_TIMEOUT;
144 static int WTMPXfd = -1; /* File descriptor of WTMPX_FILE */
145 static int WTMPX_ufreq = WTMPX_UFREQ;
146 static int Debug = 0; /* Set by command line argument */
147 static int Max_fds = MAX_FDS;

149 /*
150 * This program has three main components plus utilities and debug routines
151 * Receiver - receives the process ID or process for us to watch.
152 * (Uses a named pipe to get messages)
153 * Watcher - Use poll(2) to watch for processes to die so they
154 * can be cleaned up (get marked as DEAD_PROCESS)
155 * Scanner - periodically scans the utmpx file for stale entries
156 * or live entries that we don't know about.
157 */

159 static int wait_for_pids(); /* Watcher - uses poll */
160 static void scan_utmps(); /* Scanner, reads utmpx file */
161 static void drain_pipe(); /* Receiver - reads mesgs over UTMPPIPE */
162 static void setup_pipe(); /* For setting up receiver */

164 static void add_pid(); /* Adds a process to the table */
165 static void rem_pid(); /* Removes a process from the table */
166 static int find_pid(); /* Finds a process in the table */
167 static int proc_to_fd(); /* Takes a pid and returns an fd for its proc */
168 static void load_tables(); /* Loads up the tables the first time around */
169 static int pidcmp(); /* For sorting pids */

171 static void clean_entry(); /* Removes entry from our table and calls ... */
172 static void clean_utmpx_ent(); /* Cleans a utmpx entry */

174 static void fatal() __NORETURN; /* Prints error message and calls exit */
175 static void nonfatal(); /* Prints error message */
176 static void print_tables(); /* Prints out internal tables for Debug */
177 static int proc_is_alive(pid_t pid); /* Check if a process is alive */
178 static void warn_utmp(void);

180 /* Validate defaults from file and assign */
181 static int validate_default(char *defp, int *flag);

183 /*
184 * main() - Main does basic setup and calls wait_for_pids() to do the work
185 */

187 int
188 main(int argc, char *argv[])
189 {
190     char *defp;
191     struct rlimit rlim;
192     int i;
193     time_t curtime, now;
194     char msg[256];

196     prog_name = argv[0]; /* Save invocation name */

198     if (getuid() != 0) {
199         (void) fprintf(stderr,
200             "You must be root to run this program\n");
201         fatal("You must be root to run this program");
202     }

204     if (argc > 1) {
205         if ((argc == 2 && (int)strlen(argv[1]) >= 2) &&

```

```

206         (argv[1][0] == '-' && argv[1][1] == 'd')) {
207             Debug = 1;
208         } else {
209             (void) fprintf(stderr,
210                 "%s: Wrong number of arguments\n", prog_name);
211             (void) fprintf(stderr,
212                 "Usage: %s [-debug]\n", prog_name);
213             exit(2);
214         }
215     }

217     /*
218     * Read defaults file for poll timeout, WTMPX update frequency
219     * and maximum number of processes to monitor.
220     * Read defaults file for poll timeout
221     */
222     if (defopen(UTMP_DEFAULT) == 0) {
223         if ((defp = defread("SCAN_PERIOD=")) != NULL)
224             if (validate_default(defp, &Poll_timeout) == -1) {
225                 (void) snprintf(msg, sizeof(msg), "SCAN_PERIOD
226                     " should be a positive integer, found %s",
227                     defp);
228                 nonfatal(msg);
229             }
230         if ((defp = defread("SCAN_PERIOD=")) != NULL) {
231             Poll_timeout = atol(defp);
232             dprintf(("Poll timeout set to %d\n", Poll_timeout));
233         }

234         if ((defp = defread("WTMPX_UPDATE_FREQ=")) != NULL)
235             if (validate_default(defp, &WTMPX_ufreq) == -1) {
236                 (void) snprintf(msg, sizeof(msg),
237                     "WTMPX_UPDATE_FREQ should be a positive
238                     "integer, found %s", defp);
239                 nonfatal(msg);
240             }
241         if ((defp = defread("WTMPX_UPDATE_FREQ=")) != NULL) {
242             WTMPX_ufreq = atol(defp);
243             dprintf(("WTMPX update frequency set to %d\n",
244                 WTMPX_ufreq));
245         }
246         dprintf(("WTMPX update frequency set to %d\n", WTMPX_ufreq));

247     }

248     /*
249     * Paranoia - if polling on large number of FDs is expensive /
250     * buggy the number can be set lower in the field.
251     */
252     if ((defp = defread("MAX_FDS=")) != NULL)
253         if (validate_default(defp, &Max_fds) == -1) {
254             (void) snprintf(msg, sizeof(msg), "MAX_FDS "
255                 "should be a positive integer, found %s",
256                 defp);
257             nonfatal(msg);
258         }
259     if ((defp = defread("MAX_FDS=")) != NULL) {
260         Max_fds = atol(defp);
261         dprintf(("Max_fds set to %d\n", Max_fds));
262     }
263     dprintf(("Max fds set to %d\n", Max_fds));
264     (void) defopen((char *)NULL);

265     if (Debug == 0) {
266         /*
267         * Daemonize ourselves
268         */
269         if (fork()) {
270             exit(0);

```

```

261     }
262     (void) close(0);
263     (void) close(1);
264     (void) close(2);
265     /*
266     * We open these to avoid accidentally writing to a proc file
267     */
268     (void) open("/dev/null", O_RDONLY);
269     (void) open("/dev/null", O_WRONLY);
270     (void) open("/dev/null", O_WRONLY);
271     (void) setsid(); /* release process from tty */
272 }

274 openlog(prog_name, LOG_PID, LOG_DAEMON); /* For error messages */
275 warn_utmp(); /* check to see if utmp came back by accident */

277 /*
278 * Allocate the pidtable and fdtable. An earlier version did
279 * this as we go, but this is simpler.
280 */
281 if ((pidtable = malloc(Max_fds * sizeof (struct pidentry))) == NULL)
282     fatal("Malloc failed");
283 if ((fdtable = malloc(Max_fds * sizeof (pollfd_t))) == NULL)
284     fatal("Malloc failed");

286 /*
287 * Up the limit on FDs
288 */
289 if (getrlimit(RLIMIT_NOFILE, &rlim) == 0) {
290     rlim.rlim_cur = Max_fds + EXTRA_MARGIN + 1;
291     rlim.rlim_max = Max_fds + EXTRA_MARGIN + 1;
292     if (setrlimit(RLIMIT_NOFILE, &rlim) != 0) {
293         fatal("Out of File Descriptors");
294     }
295 } else
296     fatal("getrlimit returned failure");

298 (void) enable_extended_FILE_stdio(-1, -1);

300 if ((WTMPXfd = open(WTMPX_FILE, O_RDONLY)) < 0)
301     nonfatal("WARNING: unable to open " WTMPX_FILE " for update.");
302 nonfatal("WARNING: unable to open " WTMPX_FILE " for update.");

303 /*
304 * Loop here scanning the utmpx file and waiting for processes
305 * to terminate. Most of the activity is directed out of wait_for_pids.
306 * If wait_for_pids fails we reload the table and try again.
307 */

309 curtime = time(NULL);
310 dprintf(("utmp warning timer set to %d seconds\n", WARN_TIME));

312 for (i = 0; i < MAX_RESETS; i++) {
313     load_tables();
314     while (wait_for_pids() == 1) {
315         now = time(NULL);
316         if ((now - curtime) >= WARN_TIME) {
317             dprintf(("utmp warning timer expired\n"));
318             warn_utmp();
319             curtime = now;
320         }
321     }
322 }

324 (void) close(WTMPXfd);

```

```

326     /*
327     * We only get here if we had a bunch of resets - so give up
328     */
329     fatal("Too many resets, giving up");
330     return (1);
331 }
332 unchanged_portion_omitted

692 /*
693 *
694 */
695
696 /*
697 * add_pid - add a pid to the fd table and the pidtable.
698 *
699 *
700 */
701 static void
702 add_pid(pid_t pid)
703 {
704     int fd = 0;
705     int i = 0, move_amt;
706     int j;
707     static int first_time = 1;

709     /*
710     * Check to see if the pid is already in our table, or being passed
711     * pid zero.
712     */
713     if (pidcnt != 0 && (find_pid(pid, &j) == 1 || pid == 0))
714         return;

716     if (pidcnt >= Max_fds) {
717         if (first_time == 1) {
718             /*
719             * Print this error only once
720             */
721             nonfatal("File Descriptor limit exceeded");
722             first_time = 0;
723         }
724         return;
725     }
726     /*
727     * Open the /proc file checking if there's still a valid proc file.
728     */
729     if (pid != 0 && (fd = proc_to_fd(pid)) == -1) {
730         /*
731         * No so the process died before we got to watch for him
732         */
733         return;
734     }

736     /*
737     * We only do this code if we're not putting in the first element
738     * Which we know will be for proc zero which is used by setup_pipe
739     * for its pipe fd.
740     */
741     if (pidcnt != 0) {
742         for (i = 0; i < pidcnt; i++) {
743             if (pid <= pidtable[i].pl_pid)
744                 break;
745         }

```

```

747     /*
748     * Handle the case where we're not sticking our entry on the
749     * the end, or overwriting an existing entry.
750     */
751     if (i != pidcnt && pid != pidtable[i].pl_pid) {

753         move_amt = pidcnt - i;
754         /*
755         * Move table down
756         */
757         if (move_amt != 0) {
758             (void) memmove(&pidtable[i+1], &pidtable[i],
759                 move_amt * sizeof (struct pidentry));
760             (void) memmove(&fdtable[i+1], &fdtable[i],
761                 move_amt * sizeof (pollfd_t));
762         }
763     }
764 }

766     /*
767     * Fill in the events field for poll and copy the entry into the array
768     */
769     fdtable[i].events = 0;
770     fdtable[i].revents = 0;
771     fdtable[i].fd = fd;

773     /*
774     * Likewise, setup pid field and pointer (index) to the fdtable entry
775     */
776     pidtable[i].pl_pid = pid;

778     pidcnt++;
779     dprintf((" add_pid: pid = %d fd = %d index = %d pidcnt = %d\n",
780         (int)pid, fd, i, pidcnt));
781 }

784 /*
785 * rem_pid      - Remove an entry from the table and check to see if its
786 *               not in the utmpx file.
787 *               If i != -1 don't look up the pid, use i as index
788 *
789 * pid          - Pid of process to clean or 0 if we don't know it
790 *
791 * i            - Index into table or -1 if we need to look it up
792 *
793 * clean_it    - Clean the entry, or just remove from table?
794 */

796 static void
797 rem_pid(pid_t pid, int i, int clean_it)
798 rem_pid(pid, i, clean_it)
799     pid_t pid; /* Pid of process to clean or 0 if we don't know it */
800     int i; /* Index into table or -1 if we need to look it up */
801     int clean_it; /* Clean the entry, or just remove from table? */
802 {
803     int move_amt;

804     dprintf((" rem_pid: pid = %d i = %d", (int)pid, i));

806     /*
807     * Don't allow slot 0 in the table to be removed - utmp pipe fd
808     */
809     if ((i == -1 && pid == 0) || (i == 0)) {
810         dprintf((" - attempted to remove proc 0\n"));
811         return;

```

```

809     }

811     if (i != -1 || find_pid(pid, &i) == 1) { /* Found the entry */
812         (void) close(fdtable[i].fd); /* We're done with the fd */

814         dprintf((" fd = %d\n", fdtable[i].fd));

816         if (clean_it == CLEANIT)
817             clean_entry(i);

819         move_amt = (pidcnt - i) - 1;
820         /*
821         * Remove entries from the tables.
822         */
823         (void) memmove(&pidtable[i], &pidtable[i+1],
824             move_amt * sizeof (struct pidentry));

826         (void) memmove(&fdtable[i], &fdtable[i+1],
827             move_amt * sizeof (pollfd_t));

829         /*
830         * decrement the pid count - one less pid to worry about
831         */
832         pidcnt--;
833     }
834     if (i == -1)
835         dprintf((" - entry not found \n"));
836 }

839 /*
840 * find_pid      - Returns an index into the pidtable of the specified pid,
841 *               else -1 if not found
842 */

844 static int
845 find_pid(pid_t pid, int *i)
846 find_pid(pid, i)
847     pid_t pid;
848     int *i;
849 {
850     struct pidentry pe;
851     struct pidentry *p;

852     pe.pl_pid = pid;
853     p = bsearch(&pe, pidtable, pidcnt, sizeof (struct pidentry), pidcmp);

855     if (p == NULL)
856         return (0);
857     else {
858         *i = p - (struct pidentry *)pidtable;
859         return (1);
860     }
861 }

862 /*
863 * Pidcmp - Used by bsearch for sorting and finding process IDs.
864 */

866 static int
867 pidcmp(struct pidentry *a, struct pidentry *b)
868 pidcmp(a, b)
869     struct pidentry *a, *b;
870 {
871     if (b == NULL || a == NULL)

```

```

870         return (0);
871     return (a->pl_pid - b->pl_pid);
872 }

875 /*
876 * proc_to_fd - Take a process ID and return an open file descriptor to the
877 *             /proc file for the specified process.
878 */
879 static int
880 proc_to_fd(pid_t pid)
881 {
882     char procname[64];
883     int fd, dfd;

885     (void) sprintf(procname, "/proc/%d/psinfo", (int)pid);

887     if ((fd = open(procname, O_RDONLY)) >= 0) {
888         /*
889          * dup the fd above the low order values to assure
890          * stdio works for other fds - paranoia.
891          */
892         if (fd < EXTRA_MARGIN) {
893             dfd = fcntl(fd, F_DUPFD, EXTRA_MARGIN);
894             if (dfd > 0) {
895                 (void) close(fd);
896                 fd = dfd;
897             }
898         }
899         /*
900          * More paranoia - set the close on exec flag
901          */
902         (void) fcntl(fd, F_SETFD, 1);
903         return (fd);
904     }
905     if (errno == ENOENT)
906         return (-1);

908     if (errno == EMFILE) {
909         /*
910          * This is fatal, since libc won't be able to allocate
911          * any fds for the pututxline() routines
912          */
913         fatal("Out of file descriptors");
914     }
915     fatal(procname);          /* Only get here on error */
916     return (-1);
917 }

920 /*
921 * *** Utmpx Cleaning Utilities ***
922 */

924 /*
925 * Clean_entry - Cleans the specified entry - where i is an index
926 *             into the pid_table.
927 */
928 static void
929 clean_entry(int i)
930 {
931     struct utmpx *u;

```

```

933     if (pidcnt == 0)
934         return;

936     dprintf(("    Cleaning %d\n", (int)pidtable[i].pl_pid));

938     /*
939      * Double check if the process is dead.
940      */
941     if (proc_is_alive(pidtable[i].pl_pid)) {
942         dprintf(("    Bad attempt to clean %d\n",
943                (int)pidtable[i].pl_pid));
944         return;
945     }

947     /*
948      * Find the entry that corresponds to this pid.
949      * Do nothing if entry not found in utmpx file.
950      */
951     setutxent();
952     while ((u = getutxent()) != NULL) {
953         if (u->ut_pid == pidtable[i].pl_pid) {
954             if (u->ut_type == USER_PROCESS) {
955                 clean_utmpx_ent(u);
956             }
957         }
958     }
959     endutxent();
960 }

963 /*
964 * clean_utmpx_ent - Clean a utmpx entry
965 */

967 static void
968 clean_utmpx_ent(struct utmpx *u)
969 {
970     struct utmpx *u;
971     dprintf(("    clean_utmpx_ent: %d\n", (int)u->ut_pid));
972     u->ut_type = DEAD_PROCESS;
973     (void) time(&u->ut_xtime);
974     (void) pututxline(u);
975     updwtmpx(WTMPX_FILE, u);
976     /*
977      * XXX update wtmp for ! nonuserx entries?
978      */
979 }

_____unchanged_portion_omitted_____

1042 /*
1043 * proc_is_alive - Check to see if a process is alive AND its
1044 *               not a zombie. Returns 1 if process is alive
1045 *               and zero if it is dead or a zombie.
1046 */

1048 static int
1049 proc_is_alive(pid_t pid)
1050 {
1051     char psinfoname[64];
1052     int fd;
1053     psinfo_t psinfo;

```

```

1055     if (kill(pid, 0) != 0)
1056         return (0);          /* Kill failed - no process */

1058     /*
1059     * The process exists, so check if it's a zombie.
1060     */
1061     (void) sprintf(psinfofname, "/proc/%d/psinfo", (int)pid);

1063     if ((fd = open(psinfofname, O_RDONLY)) < 0 ||
1064         read(fd, &psinfo, sizeof(psinfo)) != sizeof(psinfo)) {
1065         /*
1066          * We either couldn't open the proc, or we did but the
1067          * read of the psinfo file failed, so pid is nonexistent.
1068          */
1069         psinfo.pr_nlwp = 0;
1070     }
1071     if (fd >= 0)
1072         (void) close(fd);

1074     /* if pr_nlwp == 0, process is a zombie */
1075     return (psinfo.pr_nlwp != 0);
1076 }

1078 /*
1079 * warn_utmp - /var/adm/utmp has been deprecated. It should no longer
1080 * be used. Applications that try to directly manipulate
1081 * it may cause problems. Since the file is no longer
1082 * shipped, if it appears on a system it's because an
1083 * old application created it. We'll have utmpd
1084 * complain about it periodically.
1085 */

1087 static void
1088 warn_utmp()
1089 {
1090     struct stat s;

1092     if (lstat(UTMP_FILE, &s) == 0 &&
1093         s.st_size % sizeof(struct utmp) == 0) {
1094         nonfatal("WARNING: /var/adm/utmp exists!\nSee "
1095                "utmp(4) for more information");
1096     }
1097 }

1099 /*
1100 * validate_default - validate and assign defaults.
1101 */

1103 static int
1104 validate_default(char *defp, int *flag)
1105 {
1106     long lval;
1107     char *endp;

1109     errno = 0;
1110     lval = strtol(defp, &endp, 10);

1112     if (errno != 0 || lval > INT_MAX || lval <= 0)
1113         return (-1);

1115     while (isspace(*endp) != 0)
1116         endp++;

1118     if (*endp != '\0')
1119         return (-1);

```

```

1121     *flag = lval;
1122     return (0);
1123 }
_____unchanged_portion_omitted_____

```

new/usr/src/cmd/utmpd/utmpd.dfl

1

118 Mon Jan 12 23:37:07 2015

new/usr/src/cmd/utmpd/utmpd.dfl

5375 utmpd(1M) core dumps when WTMPX_UPDATE_FREQ is zero

1 #

2 # Copyright 1994 Sun Microsystems, Inc. All rights reserved.

3 # Use is subject to license terms.

4 #

5 #pragma ident "%Z%M% %I% %E% SMI"

6 SCAN_PERIOD=300


```

*****
4580 Mon Jan 12 23:37:08 2015
new/usr/src/man/man1m/utmpd.1m
5375 utmpd(1M) core dumps when WTMPX_UPDATE_FREQ is zero
*****
1 \" te
2 .\" Copyright 2015 Shruti V Sampat <shrutisampat@gmail.com>
3 .\" Copyright (c) 2004, 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 UTMPD 1M \"Jan 01, 2015\"
7 .TH UTMPD 1M \"Jun 4, 2008\"
9 .SH NAME
10 utmpd \- utmpx monitoring daemon
11 .SH SYNOPSIS
12 .LP
13 .nf
14 \fButmpd\fR [\fB-debug\fR]
15 .fi

17 .SH DESCRIPTION
17 .sp
18 .LP
19 The \fButmpd\fR daemon monitors the \fB/var/adm/utmpx\fR file. See
20 \fButmpx\fR(4) (and \fButmp\fR(4) for historical information).
21 .sp
22 .LP
23 \fButmpd\fR receives requests from \fBpututxline\fR(3C) by way of a named pipe.
24 It maintains a table of processes and uses \fBpoll\fR(2) on \fB/proc\fR files
25 to detect process termination. When \fButmpd\fR detects that a process has
26 terminated, it checks that the process has removed its \fButmpx\fR entry from
27 \fB/var/adm/utmpx\fR. If the process' \fButmpx\fR entry has not been removed,
28 \fButmpd\fR removes the entry. By periodically scanning the
29 \fB/var/adm/utmpx\fR file, \fButmpd\fR also monitors processes that are not in
30 its table.
31 .SH OPTIONS
32 .sp
32 .ne 2
33 .na
34 \fB-debug\fR
35 .ad
36 .sp .6
37 .RS 4n
38 Run in debug mode, leaving the process connected to the controlling terminal.
39 Write debugging information to standard output.
40 .RE

42 .SH EXIT STATUS
44 .sp
43 .LP
44 The following exit values are returned:
45 .sp
46 .ne 2
47 .na
48 \fB0\fR
49 .ad
50 .sp .6
51 .RS 4n
52 Successful completion.
53 .RE

55 .sp
56 .ne 2
57 .na

```

```

58 \fB>\fR
59 .ad
60 .sp .6
61 .RS 4n
62 An error occurred.
63 .RE

65 .SH FILES
68 .sp
66 .ne 2
67 .na
68 \fB/etc/default/utmpd\fR
69 .ad
70 .sp .6
71 .RS 4n
72 You can set default values for the flags listed below. For example:
73 \fBSCAN_PERIOD=600\fR
74 .sp
75 The values for these flags should be greater than 0. If values read
76 from the file are found to be less than or equal to 0, or containing
77 invalid characters, the default values mentioned below are retained.
78 .sp
79 .ne 2
80 .na
81 \fBSCAN_PERIOD\fR
82 .ad
83 .sp .6
84 .RS 4n
85 The number of seconds that \fButmpd\fR sleeps between checks of \fB/proc\fR to
86 see if monitored processes are still alive. The default is 300.
87 .RE

89 .sp
90 .ne 2
91 .na
92 \fBMAX_FDS\fR
93 .ad
94 .sp .6
95 .RS 4n
96 The maximum number of processes that \fButmpd\fR attempts to monitor. The
97 default value is 4096.
98 .RE

100 .sp
101 .ne 2
102 .na
103 \fBWTMPX_UPDATE_FREQ\fR
104 .ad
105 .sp .6
106 .RS 4n
107 The number of seconds that \fButmpd\fR sleeps between read accesses of the
108 \fBwtmpx\fR file. The \fBwtmpx\fR file's last access time is used by
109 \fBinit\fR(1M) on reboot to determine when the operating system became
110 unavailable. The default is 60.
111 .RE

113 .RE

115 .sp
116 .ne 2
117 .na
118 \fB/var/adm/utmpx\fR
119 .ad
120 .sp .6
121 .RS 4n
122 File containing user and accounting information for commands such as

```

```
123 \fBwho\fR(1), \fBwrite\fR(1), and \fBlogin\fR(1).
124 .RE

126 .sp
127 .ne 2
128 .na
129 \fB\fB/proc\fR\fR
130 .ad
131 .sp .6
132 .RS 4n
133 Directory containing files for processes whose \fButmpx\fR entries are being
134 monitored.
135 .RE

137 .SH SEE ALSO
137 .sp
138 .LP
139 \fBsvcs\fR(1), \fBinit\fR(1M), \fBsvcadm\fR(1M), \fBpoll\fR(2),
140 \fBpututxline\fR(3C), \fBproc\fR(4), \fButmp\fR(4), \fButmpx\fR(4),
141 \fBattributes\fR(5), \fBsmf\fR(5)
142 .SH NOTES
143 .sp
143 .LP
144 If the filesystem holding \fB/var/adm/wtmpx\fR is mounted with options which
145 inhibit or defer access time updates, an unknown amount of error will be
146 introduced into the \fButmp\fR \fBBDOWN_TIME\fR record's timestamp in the event
147 of an uncontrolled shutdown (for example, a crash or loss of power ).
148 Controlled shutdowns will update the modify time of \fB/var/adm/wtmpx\fR, which
149 will be used on the next boot to determine when the previous shutdown occurred,
150 regardless of access time deferral or inhibition.
151 .sp
152 .LP
153 The \fButmpd\fR service is managed by the service management facility,
154 \fBsmf\fR(5), under the service identifier:
155 .sp
156 .in +2
157 .nf
158 svc:/system/utmp:default
159 .fi
160 .in -2
161 .sp

163 .sp
164 .LP
165 Administrative actions on this service, such as enabling, disabling, or
166 requesting restart, can be performed using \fBsvcadm\fR(1M). The service's
167 status can be queried using the \fBsvcs\fR(1) command.
```