OSSDEVCONF 2017

Advanced system call information tool

Kaziakhmedov Edgar
MIPT, Moscow

Outline

- A few words about strace
- System call database
- Architecture of asinfo tool
- Features
- Current state

A few words about strace

strace is a useful diagnostic, instructional, and debugging tool

man

System call database

- System call number
- Number of input arguments
- System call group:
  - File-related syscalls
  - IPC-related syscalls
  - Network-related syscalls
  - Process-related syscalls
  - Signal-related syscalls
  - File descriptor-related syscall
  - Memory mapping-related syscalls
- System call name

System call database

- Array of sysent structures

```
typedef struct sysent {
   unsigned nargs;
   int   sys_flags;
   int   sen;
   int   (*sys_func)();
   const char *sys_name;
} struct_sysent;
```

- Samples of using

```
[  0] = { 3,   TD,        SEN(read),        "read"        },
[  1] = { 3,   TD,        SEN(write),       "write"       },
[  2] = { 3,   TD|TF,     SEN(open),        "open"        },
[  3] = { 1,   TD,        SEN(close),       "close"       },
```

Architecture of asinfo tool

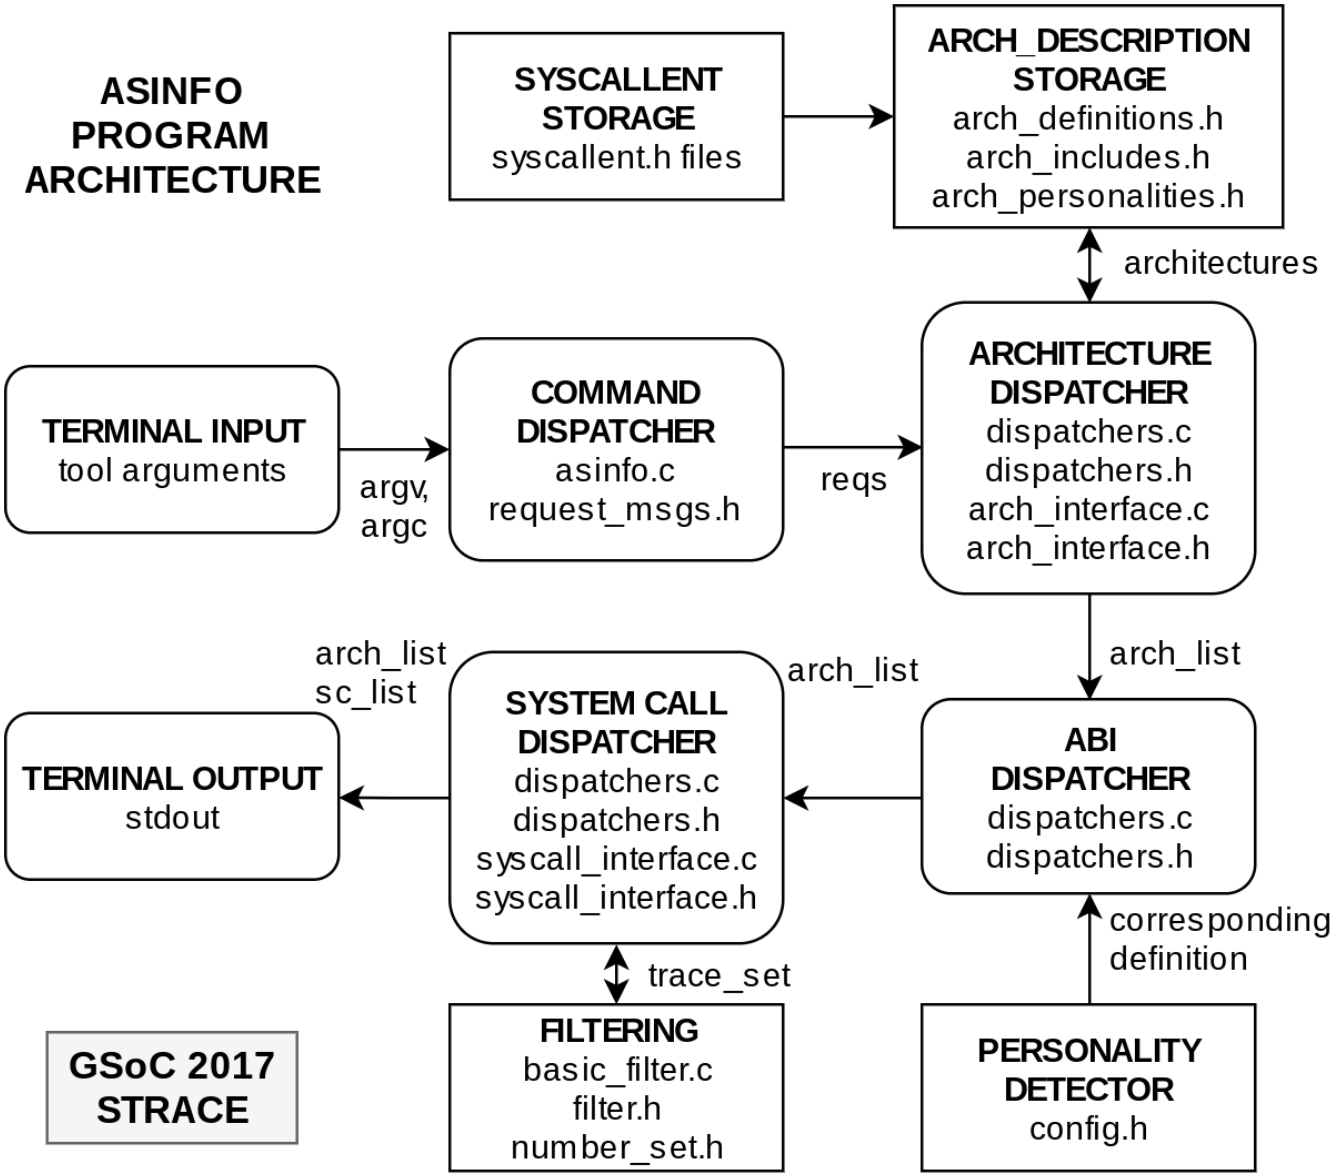

Features

- Brief information about about architecture

```
$ asinfo --set-arch or1k
| N | Architecture name | ABI mode | IMPL syscalls | IPC IMPL | SOCKET IMPL |
| 1 |     openrisc/or1k |    32bit |           277 | external |    external |
```

- Multiple architectures at a time

```
$ asinfo --set-arch aarch64,arm,blackfin,x86_64 --list-abi
| N |  Architecture name | ABI mode | IMPL syscalls | IPC IMPL | SOCKET IMPL |
| 1 |      aarch64/arm64 |    64bit |           332 | external |    external |
| 2 |      aarch64/arm64 |     eabi |           400 | external |    external |
| 3 |                arm |     oabi |           400 |  int/ext |     int/ext |
| 4 |                arm |     eabi |           400 | external |    external |
| 5 |      blackfin/bfin |    32bit |           391 | external |    external |
| 6 | x86_64/amd64/EM64T |    64bit |           333 | external |    external |
| 7 | x86_64/amd64/EM64T |      x32 |           369 | external |    external |
| 8 | x86_64/amd64/EM64T |    32bit |           381 | internal |     int/ext |
```

Features

- ABI modes

```
$ asinfo --set-arch arm7 --list-abi
| N | Architecture name | ABI mode | IMPL syscalls | IPC IMPL | SOCKET IMPL |
| 1 |               arm |     oabi |           400 |  int/ext |     int/ext |
| 2 |               arm |     eabi |           400 | external |    external |
$ asinfo --set-arch arm7 --set-abi oabi
| N | Architecture name | ABI mode | IMPL syscalls | IPC IMPL | SOCKET IMPL |
| 1 |               arm |     oabi |           400 |  int/ext |     int/ext |
$ asinfo --set-arch arm7 --set-abi oabi2
asinfo: architecture 'arm' does not have ABI mode 'oabi2'
```

- Basic operations with syscalls (strong match)

```
$ asinfo --get-sname 1
|   |              | x86_64 |
| N | Syscall name |  64bit |
| 1 |        write |      1 |
$ asinfo --get-snum write
|   |      | x86_64 |
| N | Snum |  64bit |
| 1 |    1 |  write |
```

Features

- Basic operations with syscalls (find occurrence)

```
$ asinfo --get-snum /write
|   |      |            x86_64 |
| N | Snum |             64bit |
| 1 |    1 |             write |
| 2 |   18 |          pwrite64 |
| 3 |   20 |            writev |
| 4 |  296 |           pwritev |
| 5 |  311 | process_vm_writev |
| 6 |  328 |          pwritev2 |
```

- Extended regular expression

```
$ asinfo --set-arch mips64 --set-abi o32 --get-snum "/^(.*_)?statv?fs"
|   |      |       mips64 |
| N | Snum |          o32 |
| 1 |   35 |  svr4_statfs |
| 2 |  103 | svr4_statvfs |
| 3 | 1035 |  sysv_statfs |
| 4 | 1174 | sysv_statvfs |
| 6 | 3035 | posix_statfs |
| 7 | 4099 |       statfs |
| 8 | 4255 |     statfs64 |
```

Features

- Class filtering

```
$ asinfo --get-sname ipc
|    |               | x86_64 |
|  N | Syscall name |  64bit |
|  1 |        msgctl |     71 |
|  2 |        msgget |     68 |
|  3 |        msgrcv |     70 |
|  4 |        msgsnd |     69 |
|  5 |        semctl |     66 |
|  6 |        semget |     64 |
|  7 |         semop |     65 |
|  8 |   semtimedop |    220 |
|  9 |         shmat |     30 |
| 10 |        shmctl |     31 |
| 11 |         shmdt |     67 |
| 12 |        shmget |     29 |
```

```
$ asinfo --get-sname network
|    |               | x86_64 |
|  N | Syscall name |  64bit |
|  1 |        accept |     43 |
|  2 |       accept4 |    288 |
|  3 |          bind |     49 |
|  4 |       connect |     42 |
|  5 |  getpeername |     52 |
|  6 |       getpmsg |    181 |
|  7 |  getsockname |     51 |
|  8 |   getsockopt |     55 |
|  9 |        listen |     50 |
| 10 |       putpmsg |    182 |
| 11 |     recvfrom |     45 |
| 12 |     recvmmsg |    299 |
| 13 |      recvmsg |     47 |
...
```

Features

- Discrepancies in system call characteristics

(syscall name → syscall numbers)

```
$ asinfo --set-arch amd64 --list-abi --get-sname ipc
|    |                | amd64 | amd64 | amd64 |
|  N | Syscall name   | 64bit |   x32 | 32bit |
|  1 |            ipc |     - |     - |   117 |
|  2 |         msgctl |    71 |    71 |     - |
|  3 |         msgget |    68 |    68 |     - |
|  4 |         msgrcv |    70 |    70 |     - |
|  5 |         msgsnd |    69 |    69 |     - |
|  6 |         semctl |    66 |    66 |     - |
|  7 |         semget |    64 |    64 |     - |
|  8 |          semop |    65 |    65 |     - |
|  9 |    semtimedop |   220 |   220 |     - |
| 10 |          shmat |    30 |    30 |     - |
| 11 |         shmctl |    31 |    31 |     - |
| 12 |          shmdt |    67 |    67 |     - |
| 13 |         shmget |    29 |    29 |     - |
```

Features

(syscall number → syscall names)

```
$ asinfo --set-arch amd64 --list-abi --get-snum process
|     |       |              amd64 |              amd64 |              amd64 |
|   N | Snum |              64bit |                x32 |              32bit |
|   1 |    1 |                  - |                  - |               exit |
|   2 |    2 |                  - |                  - |               fork |
|   3 |    7 |                  - |                  - |            waitpid |
|   4 |   11 |                  - |                  - |             execve |
|   5 |   56 |              clone |              clone |                  - |
|   6 |   57 |               fork |               fork |                  - |
|   7 |   58 |              vfork |              vfork |                  - |
|   8 |   59 |             execve |                  - |                  - |
|   9 |   60 |               exit |               exit |                  - |
|  10 |   61 |              wait4 |              wait4 |                  - |
|  11 |  114 |                  - |                  - |              wait4 |
|  12 |  120 |                  - |                  - |              clone |
|  13 |  158 |         arch_prctl |         arch_prctl |                  - |
|  14 |  190 |                  - |                  - |              vfork |
|  15 |  231 |         exit_group |         exit_group |                  - |
|  16 |  247 |             waitid |                  - |                  - |
...
```

Features

- System call input arguments

```
$ asinfo --set-arch amd64 --list-abi --get-sname process --nargs
|    |                    | amd64 | amd64 | amd64 |
|  N |      Syscall name  | 64bit |   x32 | 32bit |
|  1 |        arch_prctl |     2 |     2 |     2 |
|  2 |             clone |     5 |     5 |     5 |
|  3 |            execve |     3 |     3 |     3 |
|  4 |          execveat |     5 |     5 |     5 |
|  5 |              exit |     1 |     1 |     1 |
|  6 |        exit_group |     1 |     1 |     1 |
|  7 |              fork |     0 |     0 |     0 |
|  8 | rt_tgsigqueueinfo |     4 |     4 |     4 |
|  9 |           unshare |     1 |     1 |     1 |
| 10 |             vfork |     0 |     0 |     0 |
| 11 |             wait4 |     4 |     4 |     4 |
| 12 |            waitid |     5 |     5 |     5 |
| 13 |           waitpid |     - |     - |     3 |
```

Features

- Taking a guess about current architecture and ABI

X86_64, 64bit ABI

```
$ asinfo --get-sname /write
|   |                      | x86_64 |
| N |      Syscall name    |  64bit |
| 1 | process_vm_writev    |    311 |
| 2 |           pwrite64   |     18 |
| 3 |            pwritev   |    296 |
| 4 |           pwritev2   |    328 |
| 5 |              write   |      1 |
| 6 |             writev   |     20 |
```

ARM, eabi

```
$ asinfo --get-snum process
|     |       |                arm |
|   N | Snum  |               eabi |
|   1 |    1  |               exit |
|   2 |    2  |               fork |
|   3 |    7  |            waitpid |
|   4 |   11  |             execve |
|   5 |  114  |              wait4 |
|   6 |  120  |              clone |
|   7 |  190  |              vfork |
|   8 |  248  |         exit_group |
|   9 |  280  |              waitid |
|  10 |  337  |            unshare |
|  11 |  363  | rt_tgsigqueueinfo |
|  12 |  387  |           execveat |
```

Features

- Raw output

```
$ asinfo --set-arch arm64,amd64 --list-abi --get-sname /write --raw
1;pciconfig_write;-;273;-;-;-;
2;process_vm_writev;271;377;311;540;348;
3;pwrite64;68;181;18;18;181;
4;pwritev;70;362;296;535;334;
5;pwritev2;287;393;328;547;379;
6;write;64;4;1;1;4;
7;writev;66;146;20;516;146;
```

```
$ asinfo --get-sname /write --raw | awk -F ";" '{print $2,$3}' | column -t
process_vm_writev  311
pwrite64           18
pwritev            296
pwritev2           328
write              1
writev             20
```

Current state

- Merged preparatory commits
- Main commits related to asinfo is waiting for merge
- Test suites are in progress

Thank you for your attention