

# KERNEL ABI MONITORING AND TOOLCHAIN SUPPORT

`guillermo.e.martinez@oracle.com`

Linux Plumbers Conference 2022: Toolchains Track

September 14, 2022

- Interface between two binary programs.
- In the kernel:
  - Kernel $\Leftrightarrow$ Kernel.
  - Kernel $\Leftrightarrow$ Module.

# Libabigail

Libabigail is a *ABI Generic Analysis and Instrumentation* library, providing a set of tools to analyse and looks for changes in the ABI, it works with ELF binary files to:

- Extract ABI-relevant type information: functions, variables.
- Serialize & de-serialize ABI information from binary files in a XML format file.
- Compare and summarize ABI changes operating on either:
  - ELF binaries files (also specifying a path).
  - RPMs, DEBs and compress format files.
- Useful to detects API/ABI breaks in libraries.
- Verify backward compatibility with new library version.

# Info handled by Libabigail

libabigail uses as input:

- ELF objects (EXEC, DYN, REL), with debug format:
  - DWARF.
  - And now it supports: CTF!
- ELF files without debug information, just using the symtab.
- Linux Kernel image: `vmlinux` and modules: `.ko's`

# CTF: Compact C Type Format

- Associate symbols and its types.
- Lightweight debug format. So it is embedded in ELF objects.
- Debug information is organized in CTF archives and dictionaries.
  - CTF archive: `vmlinux.ctfa`.
- `-gctf` generates `.ctf` section in ELF files.
- Implementation in `libctf`.
- CTF v3: `ctf-specification`.

# Dumping CTF Info

```
1 # GNU poke/(objcopy, objdump) can be used to looks
# inside of CTF archive (vmlinux.ctfa) in a human mode.
3 #
# File vmlinux.ctfa is one of the main inputs
5 # used for CTF reader in libabigail tools.
$ objcopy --add-section .ctf=vmlinux.ctfa vmlinux vmlinux.elf
7 $ objdump --ctf --ctf-parent=shared_ctf vmlinux.elf

9 ...
CTF archive member: 3c574_cs:

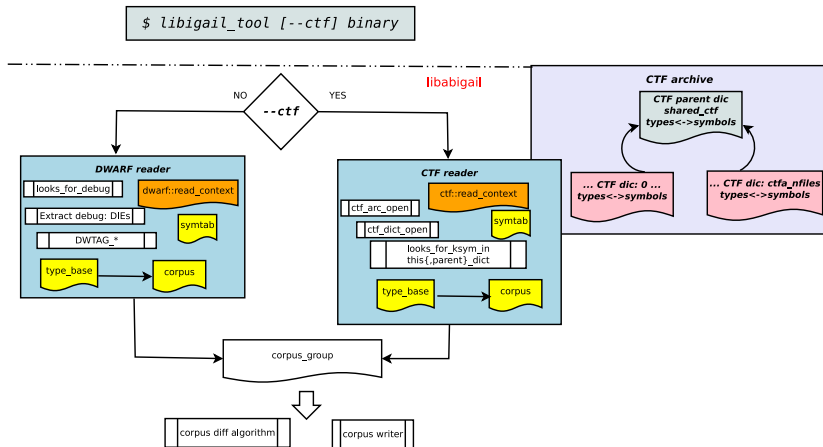
11 Header:
13   Magic number: 0xdff2
   ...
15   Parent name: shared_ctf

17 Function objects:

19 Variables:
   ...
21 _printk -> 0x855: (kind 5) int (*) (const char *, ...) (aligned at 0x8)
   ...
```

LISTING 1: dump vmlinux.ctfa

# CTF in Libabigail ...



**Figure 1:** CTF pieces used by libabigail in CTF reader.

# Libabigail Tools with CTF support

Tools with **CTF** support included in libabigail are:

- **abidw**: Represent ABI as XML output.
- **abidiff**: Summarize ABI changes between two binary files.
- **abipkgdiff**: Like abidiff but works on directories or compress/packages files.
- **kmidiff**: Compares the binary Kernel and Modules of two Linux Kernel trees.



## Example src1.c

Compare two ABIs corpus, with libabigail (diff-object), summarising absolute changes.

```
extern void other_sym();
2 const int kkkk; int a[10] = {1};
enum e
4 {
    ONE=1, TWO=2,
6 };
int
8 f0 (char p)
    { return 1; }
10 struct p
    {
12     int a;
        char b;
14 };
void
16 f2 (char c , int b, enum e e)
    { }
```

LISTING 2: src1.c

# Generating ELF with CTF debug information

```
# compile with CTF debug info
$ gcc -gctf -shared src/src1.c -o src/src1.so
```

```
$ objdump -h src/src1.so
```

```
src/src1.so:      file format elf64-x86-64
```

```
Sections:
```

Idx Name	Size	VMA	LMA	File off	Algn
...					
23 .ctf	00000168	0000000000000000	0000000000000000	00003486	2**0
	CONTENTS, READONLY				

```
# compile with DWARF and CTF debug info
$ gcc -g -gctf -shared src/src1.c -o src/src1.so
```

LISTING 3: compile.sh

# Generate ABI representation

```
$ abidw --ctf src1.so --out-file src1.abi
```

```
1 <abi-corpus version='2.1' path='src/src1.so' architecture='elf-amd-x86_64'>
  <elf-function-symbols>
3   <elf-symbol name='f0' type='func-type' binding='global-binding' ... is-defined='yes'/>
  ...
5  <elf-variable-symbols>
  <elf-symbol name='a' size='40' type='object-type' ... is-defined='yes'/>
7  ...
  <abi-instr address-size='64' language='LANG_C'>
9   <enum-decl name='e' linkage-name='e' id='type-id-3'>
  ...
11  <array-type-def dimensions='1' type-id='type-id-4' ...>
  ...
13  <qualified-type-def type-id='type-id-4' const='yes' id='type-id-8'/>
  <var-decl name='kkkk' ... mangled-name='kkkk' ... elf-symbol-id='kkkk'/>
15  <function-decl name='f0' visibility='default' ... alignment-in-bits='8' ... >
  <parameter type-id='type-id-2'/>
17  <return type-id='type-id-4'/>
  </function-decl>
19  ...
  </abi-instr>
21 </abi-corpus>
```

LISTING 4: src1.abi

## Example src2.c

```
1 extern void other_sym();
  int kkkk;
3 int a[] = {1};
  enum e
5 {
    ONE=1, TWO=2, THREE=3,
7 };
  struct p
9 {
    int a;
11 char b;
    int c;
13 };
  void
15 f0 (int p, enum e e)
    { }
17 void
  f3 (char c ,int b, struct p p)
19 { other_sym(); }
```

LISTING 5: src2.c

# Command line

```
1 # generate ABIs XML representation
$ abidw --ctf src1.so --out-file src1.abi
3 $ abidw --ctf src2.so --out-file src2.abi

5 # comparing ABIs XML references with 'libabigail' tools
$ abidiff src1.abi src2.abi

7
# comparing a reference ABI vs ELF binary file
9 $ abidiff --ctf src1.abi src2.so

11 # comparing a ELF's binary files
$ abidiff --ctf src1.so src2.so

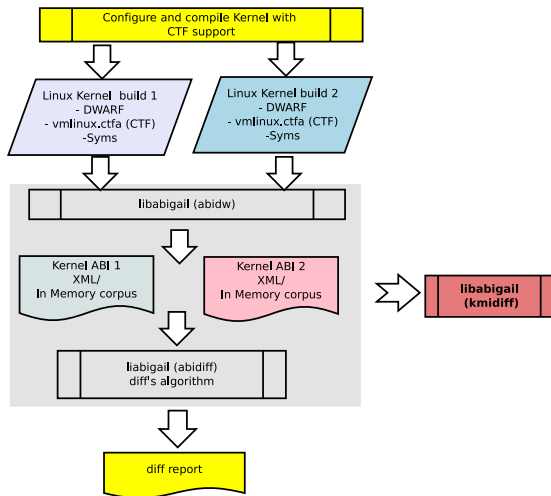
13
# comparing ABIs XML references with 'google'
15 # tools
$ stgdiff -f small -o abi.diff src1.abi src2.abi
```

LISTING 6: script1.sh

# Output report

```
Functions changes summary: 1 Removed, 1 Changed, 1 Added functions
2 Variables changes summary: 0 Removed, 2 Changed, 0 Added variables
1 Removed function:
4 [D] 'function void f2(char, int, e)' {f2}
1 Added function:
6 [A] 'function void f3(char, int, p)' {f3}
1 function with some indirect sub-type change:
8 [C] 'function int f0(char)' has some indirect sub-type changes:
  return type changed:
10   type name changed from 'int' to 'void'
   ...
12   parameter 1 of type 'char' changed:
     type name changed from 'char' to 'int'
   ...
14   parameter 2 of type 'enum e' was added
2 Changed variables:
[C] 'int a[10]' was changed to 'int a[]':
18   size of symbol changed from 40 to 4
   ...
20   type name changed from 'int[10]' to 'int[]'
[C] 'const int kkkk' was changed to 'int kkkk':
22   ...
     entity changed from 'const int' to 'int'
```

# Comparing Linux Kernel trees in libabigail



**Figure 2:** Workflow: comparing Linux kernels

# This requires CTF for Kernel (to be upstreamed)

```
1 # set CONFIG_CTF=y
  $ make && make ctf
3
  $ ls /path/to/linux_build/*
5 ..
  vmlinux
7 vmlinux.ctfa
  ...
9 module.ko
  module.ko.ctf
11 ...
```

LISTING 8: Compile Linux Kernel with CTF support



# Command line

```
1 # generate ABIs XML representation for
# Kernel + Modules
3 $ abidw --ctf --lt /path/to/build_dir_kernel1_and_modules > linux1.abi
$ abidw --ctf --lt /path/to/build_dir_kernel2_and_modules > linux2.abi
5
# comparing ABIs XML references with 'libabigail'
7 # tools
$ abidiff --impacted-interfaces linux1.abi linux2.abi > kabi.diff
9
# comparing a Kernel Linux trees
11 $ kmidiff --ctf /path/to/build_dir_kernel1_and_modules \
/path/to/build_dir_kernel2_and_modules \
13 > kabi.diff
15 # comparing ABIs XML references with 'google'
# tools
17 $ stgdiff -f small -o kabi.diff linux1.abi linux2.abi
```

LISTING 9: script2.sh

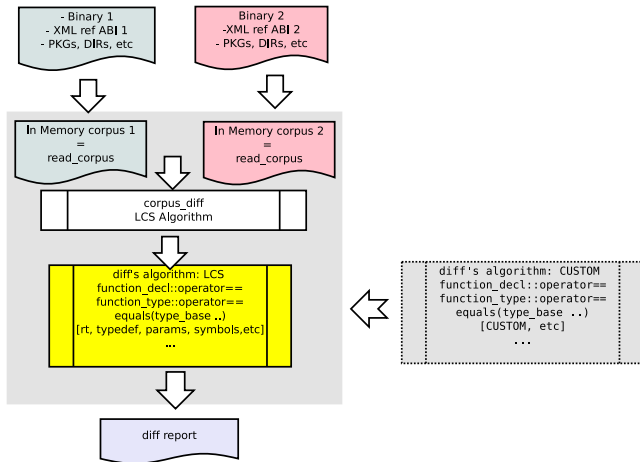
## Performance: corpus generation (CTF vs DWARF)

```
1 $ find /path/to/kernel/build -name "*.ko" | wc -l
858
3
$ ls -lh /path/to/kernel/build/vmlinux
5 -rwxr-xr-x 1 byby byby 71M Aug 18 12:18 /path/to/kernel/build/vmlinux
7
$ time abidw --ctf --lt /path/to/kernel/build > linux-ctf.abi
real    1m17.366s
9 user    0m14.778s
sys     0m1.951s
11
$ time abidw --lt /path/to/kernel/build > linux-dwarf.abi
13 real    4m47.540s
user    4m45.895s
15 sys     0m2.520s
17
$ ls -lh /path/to/kernel/build/vmlinux.ctfa
-rw-r--r-- 1 byby byby 17M Aug 14 20:10 /path/to/kernel/build/vmlinux.ctfa
19
$ objcopy --strip-debug /path/to/kernel/build/vmlinux
21 -rwxr-xr-x 1 byby byby 19M Aug 23 12:35 /path/to/kernel/build/vmlinux
```

LISTING 10: performance-script.sh

# Performance: ABI diffing

Should libabigail support custom comparison algorithms?



**Figure 3: Custom comparison algorithm.**

# Examples

```
1
2 # Percentage time reduced using libabigail tools with CTF reader.
3 #
4 -----+-----+-----+-----+-----+-----+-----+
5 | binary | version | tool | DWARF reader| CTF reader| improvement |
6 |-----+-----+-----+-----+-----+-----+-----+
7 | grep | 3.1 | abidw | 0m0.394s | 0m0.250s | 36% |
8 |-----+-----+-----+-----+-----+-----+-----+
9 | coreutils | 8.30 | abipkgdiff | 0m1.343s | 0m1.187s | 45% |
10 |-----+-----+-----+-----+-----+-----+-----+
11 | findutils | 4.6.0 | abipkgdiff | 0m0.723s | 0m0.517s | 28% |
12 |-----+-----+-----+-----+-----+-----+-----+
13 | elfutils | 0.170 | abipkgdiff | 0m0.223s | 0m0.217s | 2.6% |
14 |-----+-----+-----+-----+-----+-----+-----+
15 | vmlinux + | 5.16-rc4| abidw | 4m47.507s | 1m17.459s | 300% |
16 | modules | | | | | |
17 |-----+-----+-----+-----+-----+-----+-----+
```

LISTING 11: performance-script-us.sh

## Detecting kABI breakage: kABI v1

```
__attribute__((__noinline__))
2 const char *do_main_entry_please(const void *n, int m, int k)
{
4     printk("do_main_entry_please\n");
    if (m == k)
6         return NULL;
    else
8         return n + m;
}
10 EXPORT_SYMBOL(do_main_entry_please);

12 /* cat Module.symvers
    * 0x5e341df8 do_main_entry_please vmlinux EXPORT_SYMBOL
14 */
```

LISTING 12: kabi-v1.c

# Build a module using kABI

```
1 #include <linux/kernel.h>
  #include <linux/module.h>
3
  int hello_init(void)
5 {
    char *s = do_main_entry_please(NULL, 12, 12);
7     s = s;
    printk(KERN_INFO "Hello World!\n");
9     return 0;
  }
11
  void hello_exit(void)
13 {
    printk(KERN_INFO "Bye World!\n");
15 }
17 module_init(hello_init);
    module_exit(hello_exit);
19
    MODULE_LICENSE("GPL");
```

LISTING 13: linux-module.c

# Module's symbols

```
$ arm-linux-gnueabi-objdump -t linux-module.ko
2
linux-module.ko:      file format elf32-littlearm
4
SYMBOL TABLE:
6 00000000 l    d  .note.gnu.build-id    00000000 .note.gnu.build-id
00000000 l    d  .note.Linux      00000000 .note.Linux
8 00000000 l    d  .text 00000000 .text
00000000 l    d  .text.unlikely 00000000 .text.unlikely
10 00000000 l    d  __ksymtab      00000000 __ksymtab
00000000 l    d  __kcrctab     00000000 __kcrctab
12 0000000c l    d  .plt 00000000 .plt
...
14 00000000 g     0  .gnu.linkonce.this_module 00000240 __this_module
00000000      *UND* 00000000 __aeabi_unwind_cpp_pr0
16 00000000      *UND* 00000000 do_main_entry_please
00000018 g     F  .text.unlikely 0000000c cleanup_module
18 00000020 g     F  .text 00000028 init_module
00000000      *UND* 00000000 _printk
20 00000020 g     F  .text 00000028 hello_init
00000018 g     F  .text.unlikely 0000000c hello_exit
```

LISTING 14: module-syms.sh

```
1 __attribute__((access(read_only, 1))) __attribute__((__noinline__))
2 const char *do_main_entry_please(const void *n, int m, int k)
3 {
4     printk("do_main_entry_please\n");
5     if (m == k)
6         return NULL;
7     else
8         return n + m;
9 }
10 EXPORT_SYMBOL(do_main_entry_please);
11
12 /* cat Module.symvers
13 * 0xbce9bef2 do_main_entry_please vmlinux EXPORT_SYMBOL
14 */
```

LISTING 15: kabi-v2.c



# Loading module

```
$ insmod linux-module.ko
2 [ 20.143755] linux_module: loading out-of-tree module taints kernel.
[ 20.150282] linux_module: disagrees about version of symbol do_main_entry_please
4 [ 20.157876] linux_module: Unknown symbol do_main_entry_please (err -22)
insmod: can't insert 'linux-module.ko': invalid parameter
```

LISTING 16: kabi-load-v2.sh

# Why is the kABI incompatible?

Using: `arm-linux-gnueabi-hf-objdump -D -j.text kabi-v{1,2}.o`

```
1 00000000 <do_main_entry_please>:  
   0: b480      push {r7}  
3   2: b085      sub sp, #20  
   4: af00      add r7, sp, #0  
5   6: 60f8      str r0, [r7, #12]  
   8: 60b9      str r1, [r7, #8]  
7   a: 607a      str r2, [r7, #4]  
   c: 68ba      ldr r2, [r7, #8]  
9   e: 687b      ldr r3, [r7, #4]  
   ...  
11  1e: 4618      mov r0, r3  
   20: 3714      adds r7, #20  
13  22: 46bd      mov sp, r7  
   24: f85d 7b04  ldr.w r7, [sp], #4  
15  28: 4770      bx lr
```

LISTING 17: `kabi-vX.s`

# Why is the kABI incompatible?

Using: `abidw --ctf kabi-v{1,2}.o`

```
1 <abi-corpus version='2.1' ... architecture='elf-arm'>
  <elf-function-symbols>
3   <elf-symbol name='do_main_entry_please' type='func-type' binding='global-binding'
     visibility='default-visibility' is-defined='yes' />
  </elf-function-symbols>
5  <abi-instr address-size='64' language='LANG_C'>
    <type-decl name='char' size-in-bits='8' alignment-in-bits='8' id='type-id-1' />
7    <type-decl name='int' size-in-bits='32' alignment-in-bits='32' id='type-id-2' />
    <qualified-type-def type-id='type-id-1' const='yes' id='type-id-3' />
9    <pointer-type-def type-id='type-id-3' size-in-bits='32' ... id='type-id-4' />
    <qualified-type-def type-id='type-id-5' const='yes' id='type-id-6' />
11   <pointer-type-def type-id='type-id-6' size-in-bits='32' ... id='type-id-7' />
    <function-decl name='do_main_entry_please' visibility='default' binding='global' ...
      elf-symbol-id='do_main_entry_please'>
13     <parameter type-id='type-id-7' />
      <parameter type-id='type-id-2' />
15     <parameter type-id='type-id-2' />
      <return type-id='type-id-4' />
17   </function-decl>
    <type-decl name='void' id='type-id-5' />
19  </abi-instr>
</abi-corpus>
```

# Why is the kABI incompatible?

Using: abidiff

```
$ abidiff --ctf kabi-v{1,2}.o
```

```
2 $ echo $?
```

```
0
```

LISTING 19: abidiff-kabi.sh

# Why is the kABI incompatible?

Using: poke

```
1 load elf;
3 var fd1 = open ("kabi-v1.o", IOS_M_RDONLY);
  var fd2 = open ("kabi-v2.o", IOS_M_RDONLY);
5
  var elfv1 = Elf64_File @ fd1 : 0#B;
7 var elfv2 = Elf64_File @ fd2 : 0#B;
9 for (s in elfv1.shdr)
  {
11   var name = elfv1.get_section_name (s.sh_name);
13
14   printf "Looking differences in \"%s\"\n" name;
15   var secv1 = elfv1.get_sections_by_name (name);
16   var secv2 = elfv2.get_sections_by_name (name);
17   sdiff :a secv1[0] :b secv2[0];
17 }
19 print "Done..\n";
```

LISTING 20: cmp-elf.pk

# Why is the kABI incompatible?

```
1 $ poke -L cmp-elf.pk
  Looking differences in ""
3 Looking differences in ".text"
  Looking differences in ".data"
5 Looking differences in ".bss"
  Looking differences in ".ctf"
7 Looking differences in ".comment"
  Looking differences in ".note.GNU-stack"
9 Looking differences in ".eh_frame"
  Looking differences in ".rela.eh_frame"
11 Looking differences in ".symtab"
  Looking differences in ".strtab"
13 Looking differences in ".shstrtab"
  Done..
```

LISTING 21: run-poke.sh

# Why is the kABI incompatible?

- CRC generated by ./scripts/genksyms.
- Safely use of MODULE\_INIT\_IGNORE\_MODVERSIONS.

```
static const struct modversion_info ____versions[]
2 __used __section("__versions") = {
    { 0x47fadb66, "module_layout" },
4    { 0x92997ed8, "_printk" },
    { 0x5e341df8, "do_main_entry_please" },
6    { 0xefd6cf06, "__aeabi_unwind_cpp_pr0" },
};
```

LISTING 22: linux-module.mod.c

## Can we do better? Yes, with CTF

```
$ objdump --ctf linux-module.ko
2   ...
   Function objects:
4   ...
   _printk -> 0x30: (kind 5) int (*) (const char *, ...) (aligned at 0x8)
6   do_main_entry_please -> 0x2e: (kind 5) const char *(*) (const void *, int, int) (
   aligned at 0x8)

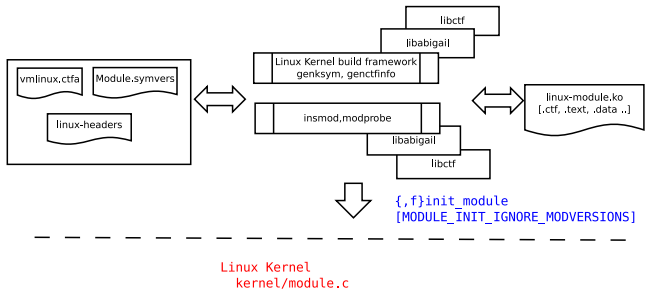
8 $ objdump --ctf --ctf-parent=shared_ctf vmlinux-v1.ctfa.elf
   ...
10  _printk -> 0x855: (kind 5) int (*) (const char *, ...) (aligned at 0x8)
   do_main_entry_please -> 0x8002ec06: (kind 5) const char *(*) (const void *, int, int) (
   aligned at 0x8)

12 $ objdump --ctf --ctf-parent=shared_ctf vmlinux-v2.ctfa.elf
14  ...
   _printk -> 0x855: (kind 5) int (*) (const char *, ...) (aligned at 0x8)
16  do_main_entry_please -> 0x8002ec06: (kind 5) const char *(*) (const void *, int, int) (
   aligned at 0x8)
```

LISTING 23: cmp-kabi-ctf.sh



# and libabigail ...



**Figure 4:** Using libabigail with CTF support.

## And why not DWARF?

- CTF size overhead 50%-100%!
- DWARF size overhead 300%!!.

```
$ ls -lh ./drivers/usb/core/usbcore.ko.ctf
2 -rw-r--r-- 1 byby byby 226K Aug 13 21:42 ./drivers/usb/core/usbcore.ko.ctf
$ cp ./drivers/usb/core/usbcore.ko ./drivers/usb/core/usbcore.ko.strip
4 $ ls -lh ./drivers/usb/core/usbcore.ko
-rw-r--r-- 1 byby byby 1.1M Aug 13 21:42 ./drivers/usb/core/usbcore.ko
6 $ objcopy -g ./drivers/usb/core/usbcore.ko.strip
$ ls -lh ./drivers/usb/core/usbcore.ko.strip
8 -rw-r--r-- 1 byby byby 293K Aug 23 15:15 ./drivers/usb/core/usbcore.ko.strip
```

LISTING 24: cmp-kabi-ctf.sh

# Unnoticed ABI breakages

```
/* gcc-12 -g -gctf abi-v1.c -c -o abi-v1.o */
2
int
4 doit(void *p, int a, int b)
{
6 char *s = (char *)p;
  if (a == b)
8   printf ("%s\n", s);
10 return a ? a : b;
}
```

LISTING 25: abi-v1.c

# Unnoticed ABI breakages

```
1 /*
   The ms_abi attribute tells the compiler to use the Microsoft ABI.
3
   x64 calling convention:
5   Integer arguments are passed in registers RCX, RDX, R8, and R9.
   Floating point arguments are passed in XMM0L, XMM1L, XMM2L, and XMM3L.
7   16-byte arguments are passed by reference...
   A scalar return value that can fit into 64 bits, including the __m64 type,
9   is returned through RAX..
   */
11
   /* gcc-12 -g -gctf ~/c/mabi.c -c -o ~/c/mabi-v2.o */
13
   __attribute__((__ms_abi__))
15 int doit(void *p, int a, int b)
   {
17   char *s = (char *)p;
   if ( a == b)
19     printf ("%s\n", s);
21   return a ? a : b;
   }
```

LISTING 26: abi-v2.c

# Unnoticed ABI breakages

```
2 $ cat test-abi.c
   extern int
4 doit(void *p, int a, int b);

6 int
   main()
8 {
   doit ("main", 12, 12);
10 }

12 gcc-12 -g -gctf -o test-abi abi-v1.o test-abi.c
   $ ./test-abi
14 main

16 gcc-12 -g -gctf -o test-abi abi-v2.o test-abi.c
   $ ./test-abi
18 Segmentation fault
```

LISTING 27: test-abi.sh

# Unnoticed ABI breakages

```
Breakpoint 1, main () at /home/byby/c/test-abi.c:5
2 5      doit ("main", 12, 12);
   (gdb) s
4 doit (p=0x555555557dd8, a=0xc, b=0xf7df3ef0) at /home/byby/c/mabi.c:4
   4      char *s = (char *)p;
6 (gdb) n
   5      if ( a == b)
8 (gdb)
   8      return a ? a : b;
10 (gdb)

12 rax      0x555555556004      0x555555556004      /* sysv arg1 */
   rbx      0x0                0x0
14 rcx      0x555555557dd8      0x555555557dd8      /* ms arg1 */
   rdx      0xc              0xc              /* ms arg2 */ /* sysv arg2 */
16 rsi      0xc              0xc              /* sysv arg3 */
   ...
18 r8       0x7ffff7df3ef0    0x7ffff7df3ef0    /* ms arg3 */
   ...
20 (gdb) p (char *) 0x555555556004
   $4 = 0x555555556004 "main"
22 (gdb)
```

LISTING 28: gdb-session.gdb

# Unnoticed ABI breakages

```
$ poke -L cmp-elf.pk
2 Looking differences in ""
Looking differences in ".text"
4 @@ 0x10+8,0x10+8 @@
-11 00 00 00 00 00 00 00 a.sh_size 0x0000000000000011UL#B
6 +12 00 00 00 00 00 00 00 b.sh_size 0x0000000000000012UL#B
Looking differences is ".data"
8 @@ 0x48+8,0x48+8 @@
-51 00 00 00 00 00 00 00 a.sh_offset 0x00000000000000051UL#B
10 +52 00 00 00 00 00 00 00 b.sh_offset 0x00000000000000052UL#B
...
12 Looking differences in ".eh_frame"
Looking differences in ".rela.eh_frame"
14 Looking differences in ".shstrtab"

16 $ abidiff abi-v{1,2}.o
$ echo $?
18 0
$ abidiff --ctf abi-v{1,2}.o
20 $ echo $?
0
```

LISTING 29: run-diff-abi.sh

# Unnoticed ABI breakages

Adding ABI metadata to ELF sections ...

- `malloc`, `nonreturn`, `section`, `zero_call_used_regs`, etc.



