

Dartmouth Computer Science Technical  
Report TR2011-680 (Draft version)  
**Exploiting the Hard-Working DWARF:  
Trojans with no Native Executable Code**

James Oakley and Sergey Bratus  
Computer Science Dept.  
Dartmouth College  
Hanover, New Hampshire  
`james.oakley.11@alum.dartmouth.org`

April 11, 2011

## Abstract

All binaries compiled by recent versions of GCC from C++ programs include complex data and dedicated code for exception handling support. The data structures describe the call stack frame layout in the DWARF format bytecode. The dedicated code includes an interpreter of this bytecode and logic to implement the call stack unwinding.

Despite being present in a large class of programs – and therefore potentially providing a huge attack surface – this mechanism is not widely known or studied. Of particular interest to us is that the exception handling mechanism provides the means for fundamentally altering the flow of a program. DWARF is designed specifically for calculating call frame addresses and register values. DWARF expressions are Turing-complete and may calculate register values based on any readable data in the address space of the process. The exception handling data is in effect an embedded program residing within every C++ process. This paper explores what can be accomplished with control of the debugging information without modifying the program's text or data. We also examine the exception handling mechanism and argue that it is rife for vulnerability finding, not least because the error states of a program are often those least well tested. We demonstrate the capabilities of this DWARF virtual machine and its suitability for building a new type of backdoor as well as other implications it has on security.

# Contents

<b>1</b>	<b>Introduction and Background</b>	<b>4</b>
1.1	Introduction . . . . .	4
1.2	Contributions . . . . .	6
1.3	History and Prior Work . . . . .	7
1.4	Exploitation at native binary code granularity . . . . .	8
1.5	Exploiting the auxiliary computations . . . . .	9
1.6	LOCREATE: a case study in programming the standard relocator. . . . .	10
<b>2</b>	<b>Technical Background</b>	<b>12</b>
2.1	Call Frame Information . . . . .	12
2.2	DWARF Expressions . . . . .	15
2.3	Exception Handlers . . . . .	16
2.4	Exception Process . . . . .	18
<b>3</b>	<b>Katana and Dwarfscript</b>	<b>19</b>
3.1	Katana Shell . . . . .	20
3.2	Dwarfscript . . . . .	21
<b>4</b>	<b>Anatomy of a DWARF Trojan</b>	<b>23</b>
4.1	Building a Dynamic Linker in DWARF . . . . .	26
4.2	Combining with Traditional Exploits . . . . .	26
<b>5</b>	<b>Limitations and Workarounds</b>	<b>29</b>
5.1	Registers and Parameter Passing . . . . .	29
5.2	No Side Effects . . . . .	30
5.3	DWARF Machine Implementation . . . . .	31
5.4	Limited <code>.eh_frame</code> space . . . . .	31

5.5	Debugging . . . . .	32
<b>6</b>	<b>Conclusion</b>	<b>33</b>
6.1	Availability . . . . .	34
6.2	Acknowledgments . . . . .	34
	<b>Bibliography</b>	<b>38</b>
<b>A</b>	<b>CIE and FDE Structure Inside .eh_frame</b>	<b>39</b>
<b>B</b>	<b>.gcc_except_table Layout</b>	<b>41</b>
<b>C</b>	<b>Dwarfsript Grammar</b>	<b>42</b>

# Chapter 1

## Introduction and Background

### 1.1 Introduction

This chapter offers a general overview of the ELF structure, its history of use in exploitation, and an overview of our contributions. Those who wish to jump into the technical details may wish to skip to Chapter 2. Historically, exploitation mostly focused on the “main” computation performed by the code of the target program and the libraries loaded into its process context (for the sake of the argument, let us define this computation as the flow described by the target program’s call graph). In ELF<sup>1</sup> terms, it was the contents of `.text` sections of executable and shared object files that received the most attention (such as being scanned for vulnerabilities, trojan logic, or “gadgets” to aid exploits).

However, a typical ELF process context is also constructed and maintained by a number of what we would call *auxiliary computations*, driven by data and/or code from other, non-`.text` ELF sections. These computations handle the special stages of the target process’ lifecycle, from creation and initialization, loading and relocation, to dynamic loading and linking in of required libraries and library functions (thus actually connecting the function-level call graph), to exception handling (the focus of this paper) and process dismantling. The sections involved include relocations sections `*.rel*`, dynamic symbol tables and other resources pointed to the directory that is the `.dynamic` segment, the pairwise related `.init`, `.fini` and `.ctors`, `.dtors`

---

<sup>1</sup>Executable and Linkable Format

sections, and so on.<sup>2</sup>

Whereas most of these auxiliary computations seem trivial, e.g., `.init` and `.fini` merely iterate over the function pointer tables in `.ctors` and `.dtors`, they may already present convenient interfaces for exploitation (e.g., [26, 17]). More complex auxiliary computation subsystems such as those responsible for dynamic linking present much richer targets [7, 33, 22], allowing for advanced techniques that co-opt the entire subsystem’s functionality, such as the pioneering PaX non-executable memory emulation and address load randomization bypass techniques [21], which, among other contributions, had co-opted the dynamic linker itself for resolution of load-time randomized addresses!

Thus exploiting auxiliary computations — those not driven by the target’s `.text` sections — has a history, on which we will comment further in Section 1.5.

We view exploitation uses of the auxiliary computations as facets of a single general phenomenon, that of programming automata responsible for these computations with data contained in the appropriate ELF sections. These sections are interpreted by their respective automata, and, when filled with crafted “program”, can leverage their computational power to accomplish a lot more than intended by their original designers and programmers, all the while not necessarily breaking the intended semantics of the automaton’s interpretation of the crafted section contents. Instead, the crafted program is abusing the additional flexibility that designers put into the automaton.

In this paper we continue this vein by demonstrating the use of the DWARF-based exception handling mechanism — the most flexible such mechanism by design, ubiquitous in modern GCC-build environments — to host a Turing-complete Trojan computation. Our Trojan “program” for the DWARF exception handling “machine” is composed entirely of valid DWARF virtual opcodes and contains no native x86 or other platform code.

One of these auxiliary computation mechanisms is the exception handling mechanism. There are several reasons why attacking exception handling is attractive. If a backdoor triggered only by exception handling is placed it may be difficult to detect. Aside from being unexpected (since this novelty will of course eventually wear off), it runs a very low risk of disrupting normal program flow. If one chooses an exception which is triggered very rarely in

---

<sup>2</sup>Besides the classic `.text`, `.data`, `.rodata`, `.bss` and others, a modern ELF file may include over 30 interrelated sections with semantically distinct functionality.

normal program operation (and indeed, in a well-engineered program any exception will fit that criterion), then there is very low risk of the backdoor being triggered/revealed unintentionally. The triggering mechanism is also already built in. When coming from the other side (i.e. not a backdoor but a direct attack on running software) exception handling is attractive because it is often not well tested. While good testing always tests the error states of a program as well as the operational states, in actual practice the error states are often tested only cursorily if at all since functionality is viewed as most important and because unlike operational states, error states are not tested upon every trial run of the software. Therefore, the likelihood of finding bugs (some of them constituting vulnerabilities) within these regions of a program is increased.

## 1.2 Contributions

We present a further step in the direction of utilizing the “auxiliary computations” to accomplish unintended, potentially malicious goals, using the arguably most powerful (in fact, Turing-complete) ubiquitous auxiliary environment to date, the DWARF exception handling mechanism.

In particular, we present a way of programming the DWARF exception-handling virtual machine environment that comes with every modern GCC-compiled exception-aware executable or shared objects files by way of providing it with crafted contents of the `.eh_frame` and `.gcc_except_table`.

We show that these contents, originally meant to flexibly and extensively accommodate present and future stack unwinding and saved register restoration logic, should be understood as powerful bytecode that allows execution of generic computations that, among other things, can read the main process memory, and in doing so make full use of the target’s dynamic symbol information.

Moreover, the bytecode is in fact very efficient at representing such symbolic memory operations, and allows us to pack much functionality into short snippets of bytecode. For example, we can package our own self-contained dynamic linker into less than 200 bytes, which easily fits within the typical `.eh_frame` section length, sparing us the effort of rebasing or relocating any other sections.

We note that this mechanism

1. involves no native executable binary code, and therefore is easily *portable*

between systems using the same or binary compatible versions of the standard exception handling libraries;

2. for the same reason, is unlikely to be checked by any current signature-based HIDS systems;
3. is *ubiquitous*, since it occurs where ever GCC-compiled C/C++ code or other exception-throwing code is supported;
4. can easily adapt to different versions of standard libraries, since the prevalence of a few standard GNU/Linux distributions means that there are only a handful of common builds, and the existing ones can be tested by universally supported code prior to deploying the attack payloads;
5. bypasses ASLR by using the system’s own facilities;
6. if combined with an appropriate memory corruption bug, can most likely be made into a non-traditional exploit payload such as a shellcode or better;
7. is unconstrained by the technical limitations of return-oriented-programming (ROP) and similar techniques that depend on a careful chaining of multiple borrowed native code gadgets: once control is given to the crafted DWARF “program” as a result of an exception, any values can be prepared and any computation can be done entirely from the DWARF virtual machine itself.

We have modified Katana, our existing academic ELF-manipulation tool to allow us to demonstrate the techniques we discuss.

### 1.3 History and Prior Work

This section sketches the history of the native binary code’s changing role in exploit programming, from straight shellcoding to “return-oriented” and other code-borrowing techniques. It then describes the alternative approaches that, unlike the latter, do not fragment the semantics of the target’s code units but rather make use of computations afforded by these units in their entirety, in effect acting as “programs” for the automata implemented by these units, not far from the implementor’s original semantics.



For example, while ROP “gadgets” are selected from the target’s loaded code — `.text` of a library, the target process, or OS kernel — without any regard for the containing unit’s semantics other techniques like [21] or [28] use original developer-intended granularity components of a loaded process image specifically for the kind of computations they were meant to provide (although not in the contexts they were meant for).

The readers already intimately familiar with these techniques are encouraged to skip to Section 2.

## 1.4 Exploitation at native binary code granularity

Historically, the prevailing notions of computer system exploitation tended to revolve around its platform’s native binary code, at the granularity of single binary or assembly instructions.

The hacker research community has long followed the paradigm of approaching exploitation as a kind of (macro) assembly programming [25, 10, 21, 24, 9] based on co-opting the binary code of the target through its (exploitable) bugs such as memory corruptions [19].

Whereas other kinds of exploitable bugs such as integer overflows, escape character (mis)interpretation, Unicode parsing ambiguities, internal command language injection (shell commands into CGI scripts, SQL injections, etc) are also recognized as valuable, the popular judgment of technical supremacy is clearly given to exploits that deliver full programmatic control of the target.

This degree of control tended to circle back to the ability to execute the binary code of attacker’s choice. Initially, this binary code existed as *shellcode*, an executable binary blob snuck into the target process’ address space by means of crafted input, and placed in the path of the CPU’s control flow in the context of the target process.<sup>3</sup>

Following the introduction of early non-executable memory countermeasures [32, 23] (i.e., prior to the adoption of the MMU-supported non-executable (NX) bit page-level tagging by the CPU/MMU vendors), the focus of ex-

---

<sup>3</sup>Kernel exploits, such as network link layer driver exploits [2], display the same basic approach, even though they must build on the kernel’s internal API or DDK when defined, rather than on the well-specified system call or standard library APIs.

exploitation research shifted to “borrowing” necessary executable code snippets from the target’s own address space [29, 37], finally achieving academic recognition as a general, Turing-complete technique through [27, 6, 14] and subsequent developments.

Still, the emphasis of this direction is on programming with *native* binary code at instruction granularity, whether injected or borrowed. In particular, successful exploitation required knowledge of, access to, and building on long chains of exact binary or assembly snippets. In contrast, if DWARF data can be injected it can directly perform computation. Instead of a target program being supplied with data to chain snippets of code painfully found in the original program, only the environment to throw an exception must be found in the target program once data injection is achieved.

## 1.5 Exploiting the auxiliary computations

However, side-by-side with instruction-granular techniques, another approach has been developing. It recognized that large units of code — spanning multiple functions and/or system calls — already present in the target by the original software engineering design could be used to perform computations of interest to the attacker by merely manipulating the input data structures to these units.<sup>4</sup>

This approach yielded rich results when applied to the omnipresent standard code that performed the auxiliary computations in the target process’ lifecycle. As those computations necessary to create, load, link, debug, handle exceptions, and finally dismantle a process got progressively more complicated with the progress of operating systems, compilers, and programming environments — the number of ELF sections typically present in a binary nearly doubled in the last decade — the subsystems that performed them got both more powerful and better defined.

As a result, many of these subsystems, for example the relocation subsystem mentioned below and the exception handling subsystem that is the subject of this paper, developed into well-defined automata with the input

---

<sup>4</sup>Arguably, this description could apply to individual system calls as used by a classical shellcode or to library calls chained by a sequence of crafted stack frames in the style of [21], if one takes into account their decomposition into internal sub-units. However, we attach significance to their authors’ idea of what constituted a natural functionality decomposition and made these units “elementary” as opposed to larger “subsystems”.

data formats what amounted to their own distinct sets of virtual instructions.

These excellent software engineering developments have been noticed and co-opted by the exploit developers. The target’s code functionality of these automata could now be borrowed not just as a matter of opportunity or convenience orthogonal to their original purpose, but rather for their original, intended function. This was in stark contrast to “return-into-lib” exploit “gadgets” and other pieces of borrowed logic that accomplished a task conducive to the exploitation such as a memory overwrite — say, as the *malloc* implementation’s doubly-linked list patching code that manages the allocated block list is commandeered to perform a memory overwrite in a “double-free” exploitation scenario [4] — and offered superior reliability of exploits and even the power to bypass protections like ASLR (e.g., [21, 10]).

## 1.6 LOCREATE: a case study in programming the standard relocater.

A great – and pedagogically perfect in its clarity — example of using a load-time auxiliary computation for a task traditionally performed by native binary code is provided by the LOCREATE unpacker [28]. LOCREATE uses the ELF process context’s own relocation mechanism as the unpacker, driving it with crafted relocation sections. In effect, it treats the relocation subsystem as a distinct memory-transforming automaton that happens to be present in the target’s context, and drives the code transformation with a crafted “program” for this automaton.

This method is in striking contrast to the traditional unpacking of injected code (packed for inconspicuous transport to avoid NIDS and HIDS signature detection), which has been traditionally handled by custom snippets of native code (themselves prone to detection). Advanced target process instrumentation techniques such as Ollybone [30] arose to catch the unpacked code just after the unpacker was done with it, right at the point when the unpacked code started executing.<sup>5</sup> LOCREATE’s proof-of-concept unpacking performed by the relocation subsystem following the supplied “relocation

---

<sup>5</sup>We regard Ollybone’s manipulation of the x86 hardware memory translation features as a first example and a harbinger of smarter memory trapping to catch composite events, in particular “page just written to was used to fetch executable code from”, expressing a two-step trappable memory event composed from a temporal, sequential relationship of elementary memory events “read, write, execute”.

program” requires no injection of native executable code whatsoever to accomplish the task.

# Chapter 2

## Technical Background

In order to understand both how the exception-handling process may be controlled to engineer an exploit and the functionality of our tool Katana it is necessary to understand how the C++ exception handling process as implemented by GCC and as partially standardized by the Linux Standards Base [1] and the x86\_64 ABI [20]. All technical details are discussed with regards to C++, GCC and Linux and with specific attention paid to the x86\_64 architecture. The concepts (and most of the details) apply equally well to other processor architectures and to the BSDs, Solaris, and most other Unix/Unix-like systems where GCC is used. In addition, the Clang C++ compiler is known to be (nearly) fully binary compatible with GCC including with largely undocumented GCC language/implementation-specific exception handler tables, as can be seen in the LLVM source [16]. Further, the exception-handling data formats and processes are not C++-specific. All study in this work has been in the context of C++ but the same general method is used for other gcc-compiled languages supporting exceptions.

### 2.1 Call Frame Information

One of the key tasks which must be undertaken during the exception handling process is that of unwinding the stack so that the exception may be handled by a handler higher up on the stack. Obviously one may walk the call stack following return address pointers to find all call frames. This is not sufficient for restoring execution, however, as it does not respect register state. Callee-saved registers will not be restored in their normal manner

when the execution path of a procedure is interrupted by an exception. It is therefore necessary that the information necessary to restore registers at the time of an unexpected procedure termination (when an exception is thrown from within the procedure) be somehow present at the time of exception throwing/handling.

This is a problem already solved for debugging, as a debugger must do a very similar task when displaying backtraces, allowing the operator to examine the local variables at various levels in the call stack, and so on. Therefore, the Call-Frame Information section of the DWARF (Debugging with Attributed Records Format) standard [11] has been adopted for encoding the unwinding information necessary for exception handling with some minor differences, particularly in the area of pointer encoding, which are for the most part documented [20] [1]. It should be noted that the current version of the DWARF standard at the time of this writing is version 4 and most of the information in this paper is drawn from that version. GCC does not in general check which version of DWARF a program was compiled against unless necessary for resolving the layout of a structure or behaviour which conflicts across standards. No checks are made when newer features are used.

Conceptually, what this unwinding information describes is a large table which for every machine instruction in the program text (the rows) describes how to restore the machine state at the previous call frame as if control were to return up the stack from that instruction. The machine state (the columns of this table) is comprised of registers and a Canonical Frame Address (CFA). DWARF allows for an arbitrary number of registers, identified merely by number. It is up to individual ABIs to define a mapping between DWARF register numbers and the hardware registers. The DWARF registers are not required to map to actual hardware registers, for example the return address is encoded as a DWARF register but will not generally correspond to a hardware register. Each cell of this table holds a rule detailing how the contents of the register will be restored for the *previous* call frame. DWARF allows for several types of rules and the curious reader is invited to find them in the DWARF standard [11]. Most registers are restored either from another register or from a memory location accessed at some offset from the CFA. The CFA is an artificial construct (i.e. internal to the DWARF encoding and interpretation) which expresses a canonical address for the call frame on the stack. Most values relevant to the execution of a procedure can therefore be found at some small offset from the CFA. An example (not taken directly from a real program, but modeled after what may be found) of a portion of

this table is given in Figure 2.1.

PC (eip)	CFA	ebp	ebx	eax	return addr.
0xf00f000	rsp+16	*(cfa-16)			*(cfa-8)
0xf00f001	rsp+16	*(cfa-16)			*(cfa-8)
0xf00f002	rbp+16	*(cfa-16)		eax=edi	*(cfa-8)
⋮	⋮	⋮	⋮	⋮	⋮
0xf00f00a	rbp+16	*(cfa-16)	*(cfa-24)	eax=edi	*(cfa-8)

Figure 2.1: Example Conceptual Unwinding Table

We note that this table if constructed in its entirety would be absurdly large. It would be considerably larger than the text of the program itself. There are, however, many empty cells and many duplicated entries in columns. Much of the DWARF call frame information standard is essentially a compression technique so that sufficient information may be provided to, at runtime, build parts of the table as needed without the full, prohibitively large, table ever being built or stored. This compression is performed by introducing the concept of Frame Description Entities (FDEs) and DWARF instructions. An FDE corresponds to a logical block of program text (often a procedure although there is no requirement on this) and describes how unwinding may be done from within that procedure. To conserve space, information common to many FDE's is separated into a Common Information Entity (CIE) which holds many of the bookkeeping details. The precise details of the CIE and FDE structures may be found in the DWARF standard. The version of the CIE structure used for `.eh_frame` derives from DWARF versions 2 and 3 and does not include new fields added to the structure in DWARF version 4. A diagrammatic view of CIE and FDE structure is shown in Appendix A. Each FDE contains a series of DWARF instructions. There are two major types of instructions. The first specifies one of the column rules (registers) as from our table above. This rule applies to all cells in that column from the *current location* to the end of the procedure unless a different rule is specified for the same column/register later in the sequence. The *current location* on which these instructions acts begins at the first program text location described by the FDE. The second type of DWARF instruction, location instructions, advances or moves the current location to which the rule instructions apply. In this manner the entire table can be specified in a much smaller form.

## 2.2 DWARF Expressions

As noted earlier, most of the register rules specify the restoration of a register from another register or from a location on the stack (relative to the CFA). DWARF was not designed for any particular hardware or software platform, however, and there was a very conscious effort to be as flexible as possible. Its designers could not anticipate all ways in which the values of registers were to be restored. Therefore, DWARF version 3 introduced the concept of DWARF expressions (they were present to a much lesser degree in DWARF version 2) which have their own set of instructions. A register may be restored as the result of a DWARF expression. A DWARF expression consists of DWARF expression operations (instructions). These operations are evaluated on a stack-machine. Most instructions operate on the top items on the stack. While the DWARF standard does not specify the data format of stack items, GCC implements them as architecture word-sized objects. All of the basic operations necessary for numerical computation are provided: pushing constant values onto the stack, arithmetic operations, bitwise operations, and stack manipulation. In addition, DWARF expressions provide instructions for dereferencing memory addresses and obtaining the values held in registers (DWARF registers calculated as part of the unwind process so far, not necessarily machine registers). This allows registers to be restored from memory locations and registers with additional arithmetic applied. This is a fairly straightforward extension of the simpler register rules provided (with the important difference that memory dereferences may be done on absolute rather than stack-relative addresses). To truly allow register restoration from arbitrarily computed values, however, DWARF expressions include conditional operations and a conditional branch instruction. Due to this extreme flexibility, there is a complete mostly unseen machine capable of arbitrary computation residing in the address space of every GCC-compiled C++ program or program linking C++ code. As an example, if a `char*` string may be found on the stack as the first local variable below the base pointer, the DWARF expression given in Listing 2.1 finds the length of that string and returns it as the result of the expression. A complete explanation of all of the instructions used can be found in the DWARF standard [11].



Listing 2.1: DWARF strlen expression

```
#value at -0x8(%rbp) on stack
DW_OP_breg6 -8
DW_OP_lit0 #initial strlen
DW_OP_swap
DW_OP_dup
#loop begins here
DW_OP_deref_size 1
#branch if top of stack nonzero
DW_OP_bra 3 #forward 3 bytes
DW_OP_skip 8 #skip to the end
#increment the counted length
DW_OP_swap
DW_OP_lit1
DW_OP_plus
DW_OP_over
#add length to char pointer
DW_OP_plus
DW_OP_skip -16 #back 16 bytes
#finally put the character
#count on the top of the stack
#as return value
DW_OP_swap
```

## 2.3 Exception Handlers

We have observed how the unwinding of stack frames and the accompanying register restoration is performed. It is also necessary to understand how exception handler (catch blocks in C++ terminology) information is encoded. DWARF is designed as a debugging format where the debugger will be in control of how far to unwind the stack. It therefore does not provide any mechanism for encoding what sort of conditions stop the unwinding process. What it does provide is the means for augmentation. Every CIE includes an augmentation string, the contents of which are implementation defined. It is designed to allow a DWARF producer (software creating the DWARF information) to communicate to a DWARF consumer (software reading the

DWARF information) which of a set of previously agreed upon augmentations to the CIE and FDE structures are being used. The augmentations to be used on Linux and x86\_64 are well-defined by the respective standards [1] [20]. The defined augmentations allow a language-specific data area (LSDA) and personality routine to be associated with every FDE. Both of these pieces of information are specified as pointers (which may be relative or absolute). When unwinding an FDE, the exception handling process is required to call the personality routine associated with the FDE. The personality routine will interpret the LSDA and determine if a handler for the exception has been found. The actual contents of the LSDA are not defined by any standard. Standardization is viewed as not necessary because the structure that is standardized allows portions of a program written in two different languages with different information that needs to be encoded about exception handlers to coexist and to even pass exceptions between each-other (to a certain degree). The language-specific portions are handled by separate personality routines for the two modules.

The result of this is that the encoding of where exception handlers are located and what type of exceptions they handle is mostly non-standardized and non-documented. What scanty documentation there is not codified in official sources but is only to be found on the gcc mailing lists in posts such as these [36],[31]. These link to an old Hewlett-Packard document [13] which describes the format in fair detail, except it is either wrong or outdated as practical experimentation will show. The best known source of information on the format used by GCC is the assembly code generated by GCC with the flags `-fverbose-asm -dA`. In an ELF binary, the section `.gcc_except_table` contains an array of LSDAs (not ordered in any particular manner, as they are reached from the LSDA pointers in augmented FDEs). Essentially, an LSDA breaks the text region described by the corresponding FDE into call sites. Call sites corresponding to code within a try block (to use C++ terminology) each have a pointer to a chain of C++ typeinfo descriptors. These objects are used by the personality routine to determine whether the thrown exception can be handled by the handler in the current frame. A diagram of LSDA structure can be found in Appendix B.

## 2.4 Exception Process

Finally, once an understanding of the data driving the exception handling process has been gained, it is important to understand the code path taken during the throwing of an exception. This path is shown in Figure 2.2




Figure 2.2: C++ Exception Code Flow

The most important facet is that `libgcc` computes the machine state as a result of the unwinding, directly restores the necessary registers, and then returns into the handler code. We note that at least in current (4.5.2) GCC implementations, this means that at the time execution is first returned to the handler code, the data from which the registers were restored will still be present below the stack pointer until it is overwritten. The handler code is known as the *landing-pad*

With this knowledge, the only barrier to building something interesting with DWARF is the lack of tools to work with it.

## Chapter 3

# Katana and Dwarfscript

DWARF is chiefly the province of compiler and debugger authors. There are several tools which allow one to examine the DWARF frame information contained within an ELF binary. Known to the authors are `readelf`, `objdump` (which uses the same `libbfd` used by `gdb`), and `dwarfdump` [3]. There are, however, no known tools which allow manipulation of DWARF structures at a high level. The author of `dwarfdump` and `libdwarf` is known to be working on a DWARF generation tool which will support, among others, a text input mode, but as of this writing he has not made available any code supporting creating DWARF objects from a textual representation. A high-level way of manipulating call frame and exception handler information is essential to examining the security implications of DWARF and demonstrating the consequences of an adversary gaining control of this information.

To bridge this gap, we present **katana**, a tool for ELF and DWARF manipulation, and Dwarfscript, a language for expressing call frame unwinding information and the corresponding exception handler information. Katana as an ELF and DWARF manipulation tool was first developed at Dartmouth College for hotpatching research [5]. The underlying binary manipulation framework was flexible enough that we were able to readily add additional exception-handling related features to it and improve its existing handling of DWARF. Katana provides an easy-to-use tool with a shell-like interface capable of emitting the Dwarfscript representation of an ELF binary. A user may modify the Dwarfscript and then use Katana to assemble it into its binary form which can be reinserted into the executable by Katana. This allows sufficient power to experiment with the consequences of carefully crafted DWARF instructions.

Katana is built on `libelf`[8] and `libdwarf`[3]. `libelf` was chosen instead of the GNU BFD because while the latter offers more features `libelf` allows very simple, bare, access to the underlying ELF structures and thus its use did not constrain the design of Katana, it merely saved labor. `libdwarf` is used when parsing binary DWARF sections. The generation of binary DWARF data is implemented natively in Katana to allow the necessary level of control and because `libdwarf` is more intended for dealing with debugging information and not designed to cope well with generating `.eh_frame`.

### 3.1 Katana Shell

Programming any automaton such as the DWARF virtual machine requires a reasonable toolchain. To aid our experimentation, it was necessary to develop two languages. The first of these is the Katana shellscript. This turns Katana into a command interpreter/shell in the spirit of Elfsh [18]. The Katana shell has not been designed to be full-featured. It is expanded as necessary for the authors research efforts, and feature requests or suggestions are welcome. It is a simple command interpreter and does not currently support any form of control flow. It supports dynamically typed variables and a small set of commands. The types of variables are as follows:

- strings
- integers
- raw data
- ELF objects
- ELF sections
- arrays

The current command-set is as follows. Some commands have return values which can be assigned to variables by the code `$VARNAME=CMD`

- `load FILENAME`  
Loads data from `FILENAME`. If it can be interpreted as an ELF object it is. Otherwise it is interpreted as a raw data object. The return value of the command is this object.

- `save OBJECT FILENAME`  
If OBJECT is an ELF object, it is finalized (necessary section headers computed, etc) and saved to FILENAME on disk. If the object is raw data it is simply written byte-for-byte to FILENAME
- `info eh OBJECT`  
Prints information about the low-level exception handling structures for the ELF object OBJECT.
- `hash elf STRING`  
Prints the elf hash (as defined in [34]) for the given STRING.
- `dwarfscript emit [SECTION_NAME] ELF FILENAME` Prints the Dwarfscript representation of the `.eh_frame` section from the given ELF object and saves it to FILENAME. Optionally, a different SECTION\_NAME can be used instead of `.eh_frame`.
- `dwarfscript compile FILENAME` Compiles the Dwarfscript file referred to by FILENAME into raw ELF binary sections. Returns an array containing ELF sections for, in order, `.eh_frame`, `.eh_frame_hdr`, `.gcc_except_table`.

## 3.2 Dwarfscript

The goal of Dwarfscript is to provide a reasonably easy means for a human to quickly make sense of the exception handling data in an ELF binary and to modify it without having to painstakingly rearrange the underlying binary structure with a hex editor. Dwarfscript attempts to encode all information found in the ELF sections `.eh_frame/.eh_frame_hdr`, which contain the DWARF unwinding information, and in `.gcc_except_table`, which contains the exception handler information. While it should be perfectly capable of representing DWARF-standard compliant debugging information as well, its focus is on supporting all exception handling information used by GCC.

Dwarfscript is to assembly as binary exception handling data is to machine code. Dwarfscript does not attempt to provide any high-level abstractions over the exception handling data but rather to present it in a form faithful to its binary structure but allowing easy readability and modification. It is a cross between a data-description language (representing the CIE, FDE,

and LSDA structures in a textual form) and an assembly language (representing DWARF instructions and expressions as an ASCII-based language). In all cases, care has been taken to follow the structure of the DWARF data as specified in the DWARF standard [11]. The lack of any high-level constructs is deliberate. Dwarfscript allows the manipulator direct control over the data structures involved. The sample of a DWARF expression shown in Listing 2.1 is a valid part of a Dwarfscript file. It is beyond the scope here to show a complete Dwarfscript file but samples may be found in the distribution of Katana and formal grammar is given in Appendix C. Familiarity with the DWARF standard should allow one to understand Dwarfscript, as the goal is to provide a textual representation of what would otherwise be encoded in a compact binary format.

To work with Dwarfscript, one is expected to use Katana to emit the Dwarfscript for the binary which is being modified. One can then edit this Dwarfscript with the changes necessary to fulfill the goal. Katana can then be used to compile the modified Dwarfscript into binary form (the standard DWARF representation) and insert it into the executable binary. To facilitate editing, the Katana distribution contains an Emacs major-mode for Dwarfscript.

The ability to extract information from a binary executable, modify the information, and insert it back into the binary is not a common one and it makes Katana quite powerful for experimenting with binary-level changes to a program.

## Chapter 4

# Anatomy of a DWARF Trojan

What might an adversary be able to do with control of the exception handling information? In the most naïve case even without complicated DWARF expressions we could redirect the flow to skip a frame when unwinding (if we know the size of the frame on the stack). One of the simplest-possible DWARF expressions allows us to simply set a register to a constant address. Using this means we can redirect any function in our target binary to “return” to any other function in our binary. By manipulating `.gcc_except_table` we can ensure that there is always a landing pad where we would like it.

To demonstrate the power of controlling the exception handling information, we discuss how the ELF binary for a simple program can be modified to yield a shell when an exception is thrown. Our example program merely takes input from `stdin` and prints a response based on the user input. The function of the program is not relevant to our purposes. What is relevant is that if the program receives an input string it is not expecting it throws an exception (the type of the exception is irrelevant here). This program, while perhaps not very interesting, certainly does nothing that would be considered dangerous. An examination of its symbol table reveals that it does not link any of the `exec` family of functions.

We modify the ELF binary for this program in such a way that it will yield a bash shell. An examination of the modified binary will not show any differences in the text or any other section which is interpreted as machine instructions or directly affects the linking of machine instructions. Especially, we do not modify the sections `.text`, `.plt`, `.got`, `.dtors`, `.dynamic`. These sections have long been known as reasonably easy ways to insert backdoors. Modifications are made only to the following ELF sections: `.eh_frame`,



```
.eh_frame_hdr, .gcc_except_table.
```

A dynamic linker is built as a DWARF expression which locates the symbol `execvpe` in `libc`. An offset is added to this address so that control will be transferred to specific suitable instructions within the function because of the difficulty of controlling parameter passing on `x86_64` as discussed in Section 5.1. The specific point that code will be transferred to in the version and build of `libc` targeted (Arch Linux `glibc 2.13-1`) is shown in Listing 4.1.

Listing 4.1: Gadget in `libc`

```
mov    %r12,%rdx
mov    %rbx,%rsi
mov    %r14,%rdi
callq  a4eb0 <execve>
```

The FDE for the function in which the exception is thrown is modified so that one of the registers is set to the result of the dynamic-linking DWARF expression. As seen in Listing 4.1 we are setting up arguments and then calling `execve`. The call we want to effectively make is `execve("/bin/bash", "/bin/bash", "-p", NULL, NULL)`. For alignment reasons, GCC typically leaves extra padding space after `.gcc_except_table` both in-memory and in the ELF file. Therefore, we have a little extra room to insert some data (which we will know the address of) after the actual LSDA data in `.gcc_except_table`. We therefore insert the data for these `execve` parameters here. We then set up the appropriate registers in Dwarfscript as shown in Listing 4.2. Obviously all addresses are specific to where the parameter data was inserted. The DWARF register number of `rbx` on `x86_64` is 3.

There is one significant problem remaining to be solved: we must somehow transfer execution to the place in `libc` we have picked. We can modify the LSDA data in `.gcc_except_table` to control where `libgcc/libstdc++` thinks handlers are located, but we cannot trivially pretend a handler exists in `libc` since we do not even know where the library will be loaded (assuming some form of library load ASLR). The solution is to use a classic return-to-`libc` attack. We take advantage of the fact that the values computed by DWARF will be temporarily placed on the stack in order to be transferred to registers immediately upon return to the handler. We therefore set the stack pointer to just below the location of the computed address in `libc` on the stack. This does introduce a dependency on particular `libgcc/libstdc++` versions for the

Listing 4.2: Dwarfscript `execve` argument setup

```
DW_CFA_val_expression r14
begin EXPRESSION
#set to address of /bin/bash
DW_OP_constu 0x400f2c
end EXPRESSION
DW_CFA_val_expression r3
begin EXPRESSION
#set to address of address of string
#array {"/bin/bash","-p",NULL}
DW_OP_constu 0x400f3a
end EXPRESSION
DW_CFA_val_expression r12
begin EXPRESSION
#set to NULL pointer
DW_OP_constu 0
end EXPRESSION
```

amount of stack space used in handling the exception to be known. One mitigation to this dependency would be to use a DWARF expression to search a small area of the stack for known values which could be used to determine an offset. We set the stack pointer (DWARF register 7 on x86\_64) to be a constant offset from the base pointer (DWARF register 6) at the time the exception is thrown as shown in Listing 4.3. We then simply modify the landing pad in the LSDA to point to a return instruction anywhere in the binary being modified. `libgcc` will transfer control to that return instruction which will return to `libc` and the process will become a shell.

Listing 4.3: Dwarfscript stack pointer setup

```
DW_CFA_val_expression r7
begin EXPRESSION
DW_OP_breg6 -96
end EXPRESSION
```

## 4.1 Building a Dynamic Linker in DWARF

Given the general-purpose computational abilities of DWARF expressions it should not be startling that we were able to build a dynamic linker in DWARF. It demonstrates the power of DWARF used in exploits, however and shows another way that address-space layout randomization (ASLR) can be defeated. The only assumptions made are that the `.dynamic` section will not be moved by the loader and that the order that shared libraries are loaded will be the order that they are listed in `.dynamic`. This second assumption is not crucial and the dynamic linker code could easily have been slightly expanded to search for the `libc` linkmap entry. It is important to note that our DWARF dynamic linker does not simply call the standard linker in `ld.so`. It traverses the linkmap, hash-table and chain structures directly and thus is not affected by any protections built into the standard linker, such as protection against calling `dlsym` from arbitrary locations. It serves no purpose to discuss the precise details of a dynamic linker here: it is well documented elsewhere [34, 15], and ours is not substantially different in structure. Writing a dynamic linker in DWARF is similar to writing a dynamic linker in assembly except that there are no registers available: there is only a stack instead. A brief outline of the procedure is shown in Listing 4.4 to highlight what sorts of operations DWARF expressions are capable of. Those interested in the actual DWARF code are invited to contact the authors directly, and it will soon be made available in the Katana distribution.

## 4.2 Combining with Traditional Exploits

What we have concretely demonstrated so far is a trojan technique. The potential of DWARF-based exploits extends beyond this, however. If we have a means of overwriting exception handling data or otherwise making data we control to be interpreted as exception handling data, DWARF can be used in the construction of an exploit delivered at runtime (rather than a trojan).

Hypothetically, a data-injection exploit which may be unable to insert directly executable code could be used to inject DWARF bytecode which will be executed when an exception is thrown. If done, this would aid in circumventing non-executable stacks and heaps and may be an alternative to return-oriented programming requiring less careful piecemeal construc-

Listing 4.4: Dwarfscript Dynamic Linker Pseudocode

```
dereference DT_PLTGOT+8
top of stack contains a link_map*
while top of stack is not libc linkmap
    top of stack=linkmap->l_next
top of stack=linkmap->l_addr
find DT_HASH, DT_STRTAB, DT_SYMTAB
index into hashtable by "execvpe" hash
while execvpe symbol not found
    compare symbol name and "execvpe"
    if not equal
        symbol found=next in chain
get st_value from symbol
add offset to desired entry point
```

tion and with fewer known detection/mitigation techniques. If this is to be achieved, however, there must be some way to either overwrite `.eh_frame` or to fool the exception handling mechanism into reading data supplied by an attacker instead of the original `.eh_frame`.

There exist many C++ libraries in the wild with `.eh_frame` sections which are loaded read-write. Until 2002 all `.eh_frame` sections were read-write. In 2002 GCC began emitting read-only `.eh_frame` sections unless relocations were necessary for `.eh_frame` [12]. This meant that most PIC code (i.e. libraries) still required writable `.eh_frame`. Modern versions of GCC are now capable of emitting `.eh_frame` sections that do not require relocation even in PIC code, but on up-to-date distributions it is still possible to find libraries with writable `.eh_frame` sections, notably several distributions of the JVM. This is an avenue down which our future work will travel.

It would be even more beneficial to insert an alternate `.eh_frame`, since as time progresses finding binaries with writable exception handling information will become increasingly uncommon, and even at present only a fairly small set of programs will be vulnerable to `.eh_frame` overwriting. `libgcc` locates the `.eh_frame` section through use of the `GNU_EH_FRAME` program header, which points to `.eh_frame_hdr`, which in turn points to `.eh_frame`. Therefore, this program header controls where the DWARF data is read from. Overwriting program headers at runtime is not generally feasible. `libgcc`

obtains this program header through the function `dl_iterate_phdr` which is located in `libld`. Again, not an easy target. `libgcc` caches program headers, however, in the variable `frame_hdr_cache`, which is static to its compilation unit. As this is a static variable its symbol is not exported, and thus its precise location will depend on the build. While `libgcc` exports no symbols, after an exception returns there will be text addresses within `libgcc` on the invalid portion of the stack, some up to nearly 1k below the stack pointer. It is likely possible to correlate these addresses with the data address of the header cache. We have confirmed that overwriting this header cache can allow the interpretation of arbitrary data as `.eh_frame`. Fully weaponizing this into a workable exploit remains as future work.

# Chapter 5

## Limitations and Workarounds

There are several limits which anyone attempting to use the DWARF virtual machine for arbitrary computation will encounter.

### 5.1 Registers and Parameter Passing

First, it is important to note that not all machine registers may be restored during stack unwinding. The hardware ABI will define some set of registers which are callee-saved and some set which are not guaranteed to be saved. It is reasonable to assume that the unwinding implementation will restore all callee-saved registers as specified in the DWARF instructions. We have determined, both empirically and through examination of the GCC source code, that at least on the x86\_64 platform, the values assigned to certain registers through DWARF will be ignored when restoring the call frame to return into. Some of the non callee-saved registers are used for specific exception-handling purposes (i.e. on x86\_64, `rdx` is used to store an identifier for the type of exception thrown) and some appear simply not to be restored to any value. This presents a problem on architectures like x86\_64 where registers are used for passing function parameters. The authors of a backdoor would like to be able to return execution to the beginning of some interesting function (such as `execv` in `libc`) and pass it crafted arguments. The x86\_64 registers used for argument passing (`rdi`, `rsi`, `rdx`, `rcx`, `r8`, `r9`) cannot be restored, however, making this impossible. It therefore becomes necessary to return to a point inside the target function which makes use of registers that can be controlled. If the target function resides in a library this

can make it difficult for an exploit to be portable across multiple versions or even builds of the library. This problem can be mitigated to some degree if it is possible to find a suitable landing pad in the binary being modified which calls a function pointer. If the address of the library function can be called as the function pointer then dependence on the precise build of library functions will be reduced. This problem is also lessened on architectures such as 32-bit x86 which primarily use the stack rather than registers for argument passing. Another workaround for this difficulty is to code the exploit for several versions/builds. As long as it is possible to make a value appear on the target's stack (this can be in innocuous means, such as expected user input) this value can be searched for the DWARF program and used as a parameter, allowing the code path taken to be adjusted when the attack is triggered rather than merely at the time of injection.

## 5.2 No Side Effects

Through control of the stack and base pointers, a DWARF program can to some degree control the contents of the stack when execution is resumed. DWARF instructions/expressions do not, however, have the ability to directly modify memory or push anything onto the stack therefore it can be difficult to make the code at the landing pad access values on the stack that were calculated by the DWARF expression. One workaround to this is to exploit the fact that values computed for machine registers will still be in memory (at least until overwritten by new stack frames) as they are computed in memory and then transferred into the correct registers. The DWARF program can set the stack or base pointer to point to the correct region of memory. There are two limitations to this technique. The first is that the appropriate location where the computed values will be found is highly dependent on the precise stack layout and thus varies between versions and builds of libgcc and libstdc++. The second limitation is that there is of course limited contiguous stack space that can be controlled, limited by the number of restored registers and how they are laid out in memory.

## 5.3 DWARF Machine Implementation

Obviously the implementation of the DWARF virtual machine has some effect upon what sort of computations can be performed. The current (gcc 4.5.2) implementation in libgcc allows the DWARF stack to grow only to a size of 64 words [12]. The DWARF standard does not specify the maximum size of the stack and there does not appear to be a reasoned process behind this number, rather it appears to have been arbitrarily chosen as a size which should be large enough for any DWARF program gcc would expect to be necessary. While this limit should be kept in mind when writing a DWARF program it does not seriously hamper the creation of interesting DWARF programs. As discussed in Section 4.1, a dynamic linker can be programmed in DWARF using only 16 words on the stack.

## 5.4 Limited `.eh_frame` space

When modifying an ELF binary, we cannot count on the presence of full relocation information as GCC/ld does not by default emit relocatable ELF objects. Therefore, we definitely cannot count on being able to expand `.eh_frame` and we would like to avoid moving it as well. Therefore, we must be careful that DWARF programs and other modifications to FDE, CIE, and LSDA structure do not require expanding the size of `.eh_frame`. This limitation can be fairly easily overcome, however. GCC does not attempt to perform any static analysis to determine whether the call frame for a given function will ever be unwound during exception handling. `.eh_frame` will even be generated for C compilation units despite the fact that C does not support exception handling. The reason for this is that it allows exceptions to propagate seamlessly across areas of code which do not know how to deal with them. Human analysis of the program being modified, however, should yield insight into finding FDEs corresponding to functions which will never need to be unwound. In Dwarfscript, these FDEs can simply be removed from the script file, making available more room for lengthy DWARF expressions. The dynamic linker referenced earlier and discussed in Section 4.1 requires less than 200 bytes of space for its instructions.



## 5.5 Debugging

There are presently no tools available to debug DWARF programs. Rudi-mentary debugging can be achieved by stepping through the execution of the DWARF virtual machine in libgcc with a debugger. DWARF debugging support is a planned feature for Katana.

# Chapter 6

## Conclusion

We have demonstrated how the hitherto largely unexplored DWARF-format exception handling information used on a wide-variety of Unix and Unix-like platforms can be used to control the flow of execution. This has several advantages over traditional backdoors and over return-oriented-programming. Advantages of our technique include the following

- Turing-complete environment.
- ASLR defeating.
- Less likely to be detected by traditional executable-content scanners.
- Built-in trigger mechanism (the attack can lie dormant until an exception is thrown).
- Fewer carefully chained gadgets required in the target program than in return-oriented-programming therefore less analysis and time may be necessary to develop an attack.
- Does not rely on bugs. Our DWARF programs leverage existing mechanisms as an extension of their intended purpose and do not rely on implementation bugs and outright security holes but on deliberately made behavior and mechanisms.

Notably, DWARF expressions can read registers and process memory, and, as we have shown, they are suitable for writing a fully functional dynamic linker.

We stress the security risks associated with powerful computational environments added in unexpected places. While being a sterling example of extensible software engineering and introducing a conceptually graceful method of handling complex datastructures of previously unprecedented complexity, this DWARF subsystem also far exceeds expectations of most developers and defenders regarding the computations it is capable of performing.

This may lead defenders to underestimate the risks posed by such environments, and miss a number of possible attack vectors.

Finally, we release Katana as a tool to painlessly create and experiment with the sort of crafted DWARF programs we have discussed, so that interested researchers can further explore the relevant attack surface.

## 6.1 Availability

Katana is available under the GNU General Public License and may be found at <http://katana.nongnu.org/>.

## 6.2 Acknowledgments

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

This material is based upon work supported by the National Science Foundation under Grant No. CNS-0524695. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the

author(s) and do not necessarily reflect the views of the National Science Foundation.

We owe a debt of thanks to projects and people who have inspired us or have helped us along the way. These include (but are not necessarily limited to) Julien Vanegue (the ERESI project [35]), The Grugq (*Cheating the ELF* [33]), Nergal (*The advanced return-into-lib(c) exploits: PaX case study* [21]), Skape (*LOCREATE* [28]), and Tim Yardley (for helpful advice). We are of course grateful to Professor Sean Smith and the Dartmouth College Trust Lab for continuing support.

# Bibliography

- [1] Linux standard base core specification 4.0. [http://refspecs.linux-foundation.org/LSB\\_4.0.0/LSB-Core-generic/](http://refspecs.linux-foundation.org/LSB_4.0.0/LSB-Core-generic/).
- [2] The month of kernel bugs archive. <http://projects.info-pull.com/mokb/>, Nov. 2006.
- [3] ANDERSON, D. Libdwarf and dwarfdump. <http://reality.sgiweb.org/davea/dwarf.html>, Jan 2011.
- [4] ANONYMOUS. Once upon a free(). *Phrack Magazine* 57, 9 (Aug 2001).
- [5] BRATUS, S., OAKLEY, J., RAMASWAMY, A., SMITH, S., AND LOCASTO, M. Katana: Towards patching as a runtime part of the compiler-linker-loader toolchain. *International Journal of Secure Software Engineering (IJSSE)* 1, 3 (2010), 1–17.
- [6] BUCHANAN, E., ROEMER, R., SHACHAM, H., AND SAVAGE, S. When good instructions go bad: generalizing return-oriented programming to RISC. In *ACM Conference on Computer and Communications Security* (2008), pp. 27–38.
- [7] CESARE, S. Shared library redirection via ELF PLT infection. *Phrack Magazine* 56, 7 (May 2000).
- [8] DREPPER, U., MCGRATH, R., AND MACHATA, P. libelf. <https://fedorahosted.org/elfutils/>.
- [9] DULLIEN, T., KORNAU, T., AND WEINMANN, R.-P. A framework for automated architecture-independent gadget search. WOOT '10: 4th USENIX Workshop on Offensive Technologies.
- [10] DURDEN, T. Bypassing pax aslr protection. *Phrack Magazine* 59, 9 (Jul 2002).
- [11] DWARF DEBUGGING INFORMATION FORMAT COMMITTEE. DWARF debugging information format version 4. <http://dwarfstd.org/>, June 2010.
- [12] FREE SOFTWARE FOUNDATION. GCC 4.5.2 source code. <http://gcc.gnu.org/releases.html>, December 2010.
- [13] HEWLETT PACKARD CORPORATION. aC++ A.01.15. <http://www.codesourcery.com/public/cxx-abi/exceptions.pdf>, Dec 2001.

- [14] HUND, R., HOLZ, T., AND FREILING, F. Return-oriented rootks: Bypassing kernel code integrity protection mechanisms. USENIX Association.
- [15] LEVINE, J. R. *Linkers and Loaders*. Morgan-Kaufman, 1999.
- [16] LLVM TEAM. DwarfException.cpp. <http://llvm.org/svn/llvm-project/llvm/trunk/lib/CodeGen/AsmPrinter/DwarfException.cpp>. Evidence of llvm support of DWARF and gcc\_except\_table format.
- [17] LZIK. Abusing .ctors and .dtors for fun 'n profit.
- [18] MAYHEM@DEVHELL.ORG. Embedded elf debugging : the middle head of cerberus. *Phrack Magazine 11*, 61 (2003).
- [19] MEER, H. Memory corruption attacks: The (almost) complete history. Black Hat.
- [20] MICHAEL MATZ, JAN HUBI KA, A. J. M. M. *System V Application Binary Interface: AMD64 Architecture Processor Supplement*, draft version 0.99.5 ed., September 2010.
- [21] NERGAL. The advanced return-into-lib(c) exploits: Pax case study. *Phrack Magazine 58*, 4 (Dec 2001).
- [22] O'NEILL, R. Modern day elf runtime infection via got poisoning.
- [23] OPENWALL PROJECT. Linux kernel patch from the openwall project. <http://www.openwall.com/linux/>.
- [24] RICHARTE, G. About exploits writing. [http://corelabs.coresecurity.com/index.php?module=Wiki&action=attachment&type=area&page=Vulnerability\\_Research&file=publication%2FAbout\\_Exploits\\_Writing%2F2002.gera>About\\_Exploits\\_Writing.pdf](http://corelabs.coresecurity.com/index.php?module=Wiki&action=attachment&type=area&page=Vulnerability_Research&file=publication%2FAbout_Exploits_Writing%2F2002.gera>About_Exploits_Writing.pdf), 2002.
- [25] RICHARTE, G., AND RIQ. Advances in format string exploitation. *Phrack Magazine 59*, 7 (July 2002).
- [26] RIVAS, J. M. B. Overwriting the .dtors section.
- [27] SHACHAM, H. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security* (New York, NY, USA, 2007), CCS '07, ACM, pp. 552–561.
- [28] SKAPE. Lcreate: an anagram for relocate. *Uninformed 6* (Jan 2007).
- [29] SOLARDESIGNER. Getting around non-executable stack (and fix). Bugtraq mailing list, <http://www.securityfocus.com/archive/1/7480>, Aug. 1997.
- [30] STEWART, J. ollybone: Semi-automatic unpacking on ia-32.
- [31] TAYLOR, I. L. Re: Elf-section .gcc\_except\_table. <http://gcc.gnu.org/ml/gcc-help/2010-09/msg00116.html>, September 2010. gcc-help@gcc.org Mailing List.
- [32] TEAM, P. Pax project. <http://pax.grsecurity.net/>.
- [33] THE GRUGQ. Cheating the elf: subversive dynamic linking to libraries.


- [34] TIS COMMITTEE. *elf*, 1995.
- [35] VANEGUE, J., MEDEIROS, J. A. D., BISOLFATI, E., DESNOS, A., FIGUEREDO, T., GARNIER, T., LESNIAK, R., PALENCIA, J., ROY, S., SOUDAN, S., WOLOSZYN, M., AND ZABROCKI, A. The ERESI reverse engineering software interface. <http://www.eresi-project.org/>, 2009.
- [36] VARIOUS. gcc@gcc.org mailing list: Exception handling description. <http://gcc.gnu.org/ml/gcc/2009-05/msg00390.html>, May 2009.
- [37] WOJTCZUK, R. Deafeating solar designer's non-executable stack patch. <http://insecure.org/spl0its/non-executable.stack.problems.html>, January 1998.





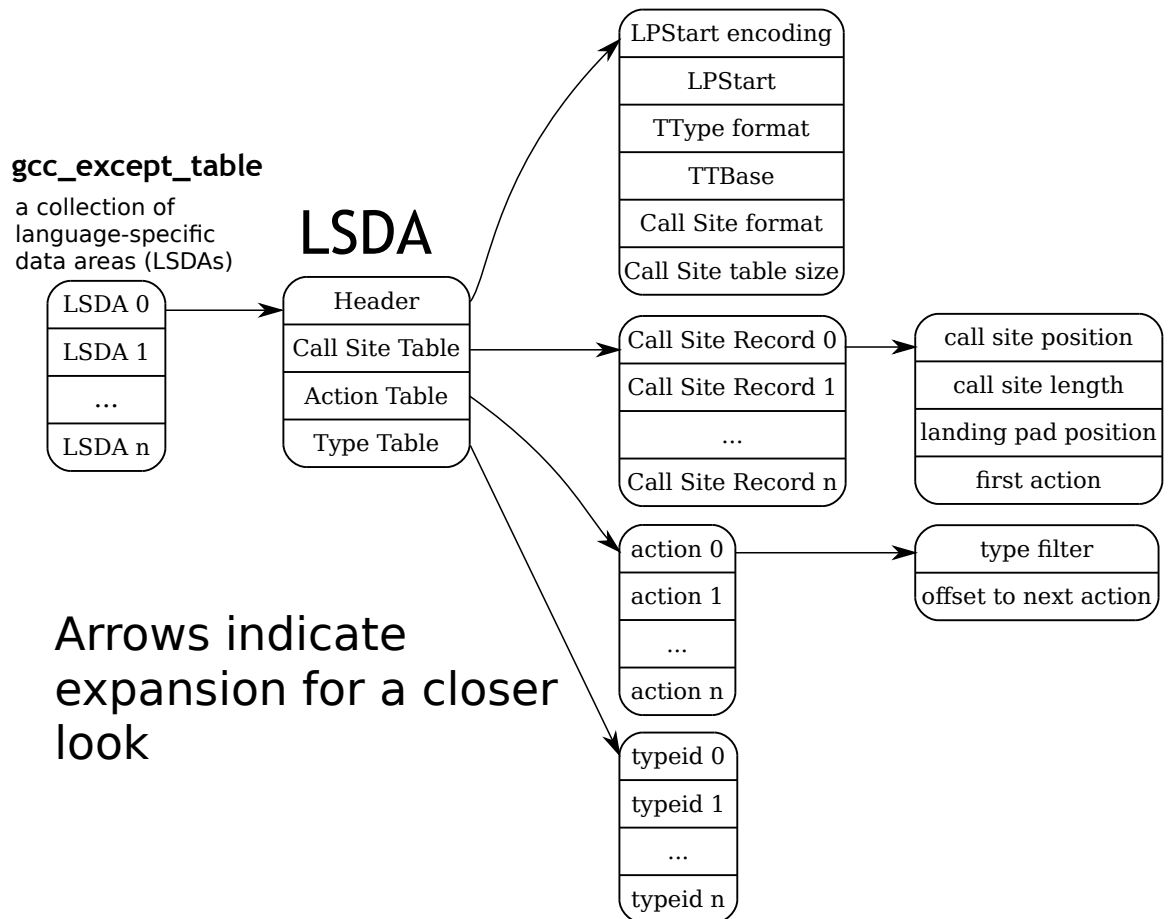
# Appendix A

## CIE and FDE Structure Inside .eh\_frame



# Appendix B

## .gcc\_except\_table Layout



# Appendix C

## Dwarfsript Grammar

`<dwarfsript> → <top_property_stmt_list> <section_list>`

`<section_list> → <section_list> <fde_section>`

`| <section_list> <cie_section>`

`| <section_list> <lsda_section>`

`| ε`

`<fde_section> → begin fde <fde_property_stmt_list> <instruction_section>`

`<fde_property_stmt_list> end fde`

`<cie_section> → begin cie <cie_property_stmt_list> <instruction_section>`

`<cie_property_stmt_list> end cie`

`<instruction_section> → begin instructions <instruction_stmt_list>`

`end instructions`

`<expression_section> → begin dwarf_expr <expr_stmt_list>`

`end dwarf_expr`

`<lsda_section> → begin lsda <lsda_part_list> end lsda`

`<call_site_section> → begin call_site <call_site_property_stmt_list>`

`end call_site`

$\langle \text{action\_section} \rangle \rightarrow \mathbf{begin\ action} \langle \text{action\_property\_stmt\_list} \rangle$   
 $\mathbf{end\ action}$

$\langle \text{top\_property\_stmt\_list} \rangle \rightarrow \langle \text{top\_property\_stmt\_list} \rangle \langle \text{top\_property\_stmt} \rangle$   
 $\langle \text{fde\_property\_stmt\_list} \rangle \rightarrow \langle \text{fde\_property\_stmt\_list} \rangle \langle \text{fde\_property\_stmt} \rangle$

$\langle \text{cie\_property\_stmt\_list} \rangle \rightarrow \langle \text{cie\_property\_stmt\_list} \rangle \langle \text{cie\_property\_stmt} \rangle$

$\langle \text{instruction\_stmt\_list} \rangle \rightarrow \langle \text{instruction\_stmt\_list} \rangle \langle \text{instruction\_stmt} \rangle$   
 $| \epsilon$

$\langle \text{expr\_stmt\_list} \rangle \rightarrow \langle \text{expr\_stmt\_list} \rangle \langle \text{expr\_stmt} \rangle$   
 $| \langle \text{expr\_stmt\_list} \rangle \langle \text{label} \rangle$   
 $| \epsilon$

$\langle \text{lsda\_part\_list} \rangle \rightarrow \langle \text{lsda\_part\_list} \rangle \langle \text{lsda\_part} \rangle$   
 $| \epsilon$

$\langle \text{call\_site\_property\_stmt\_list} \rangle \rightarrow \langle \text{call\_site\_property\_stmt\_list} \rangle$   
 $\langle \text{call\_site\_property\_stmt} \rangle$   
 $| \epsilon$

$\langle \text{action\_property\_stmt\_list} \rangle \rightarrow \langle \text{action\_property\_stmt\_list} \rangle \langle \text{action\_property\_stmt} \rangle$   
 $| \epsilon$

$\langle \text{top\_property\_stmt} \rangle \rightarrow \langle \text{section\_type\_prop} \rangle$   
 $| \langle \text{section\_location\_prop} \rangle$   
 $| \langle \text{eh\_hdr\_location\_prop} \rangle$   
 $| \langle \text{except\_table\_addr\_prop} \rangle$

$\langle \text{cie\_property\_stmt} \rangle \rightarrow \langle \text{index\_prop} \rangle$   
 $| \langle \text{length\_prop} \rangle$   
 $| \langle \text{version\_prop} \rangle$   
 $| \langle \text{augmentation\_prop} \rangle$

|<fde\_ptr\_enc\_prop>  
| <fde\_lsda\_ptr\_enc\_prop>  
|<personality\_ptr\_enc\_prop>  
| <personality\_prop>  
|<address\_size\_prop>  
| <segment\_size\_prop>  
|<data\_align\_prop>  
| <code\_align\_prop>  
| <return\_addr\_rule\_prop>

fde\_property\_stmt →<index\_prop>  
| <length\_prop>  
| <cie\_index\_prop>  
| <initial\_location\_prop>  
| <address\_range\_prop>  
| <lsda\_idx\_prop>

<lsda\_part> → <lsda\_property\_stmt>  
| <call\_site\_section>  
| <action\_section>

<lsda\_property\_stmt> → <lpstart\_prop>  
| <typeinfo\_enc\_prop>  
| <typeinfo\_prop>

<call\_site\_property\_stmt> → <position\_prop>  
| <length\_prop>  
| <landing\_pad\_prop>  
| <has\_action\_prop>  
| <first\_action\_prop>

<action\_property\_stmt> → <type\_idx\_prop>  
|<next\_prop>

<index\_prop> → **index** : <nonneg\_int\_lit>

<length\_prop> → **length** : <nonneg\_int\_lit>

<cie\_index\_prop> → **cie\_index** : <nonneg\_int\_lit>

<initial\_location\_prop> → **initial\_location** : <nonneg\_int\_lit>

<address\_range\_prop> → **address\_range** : <nonneg\_int\_lit>

<version\_prop> → **version** : <nonneg\_int\_lit>

<fde\_ptr\_enc\_prop> → **fde\_ptr\_enc** : <dw\_pe\_lit>

<fde\_lsda\_ptr\_enc\_prop> → **fde\_lsda\_ptr\_enc** : <dw\_pe\_lit>

<personality\_ptr\_enc\_prop> → **personality\_ptr\_enc** : <dw\_pe\_lit>

<personality\_prop> → **personality** : <nonneg\_int\_lit>

<address\_size\_prop> → **address\_size** : <nonneg\_int\_lit>

<segment\_size\_prop> → **segment\_size** : <nonneg\_int\_lit>

<data\_align\_prop> → **data\_align** : <int\_lit>

<code\_align\_prop> → **code\_align** : <nonneg\_int\_lit>

<return\_addr\_rule\_prop> → **ret\_addr\_rule** : <nonneg\_int\_lit>

<section\_type\_prop> → **section\_type** : <string\_lit>

<section\_location\_prop> → **section\_loc** : <nonneg\_int\_lit>

<eh\_hdr\_location\_prop> → **eh\_hdr\_loc** : <nonneg\_int\_lit>  
 <except\_table\_addr\_prop> → **except\_table\_addr** : <nonneg\_int\_lit>  
 <lpstart\_prop> → **lpstart** : <nonneg\_int\_lit>  
 <typeinfo\_enc\_prop> → **typeinfo\_enc** : <dw\_pe\_lit>  
 <typeinfo\_prop> → **typeinfo** : <nonneg\_int\_lit>  
 <position\_prop> → **position** : <nonneg\_int\_lit>  
 <landing\_pad\_prop> → **landing\_pad** : <nonneg\_int\_lit>  
 <has\_action\_prop> → **has\_action** : <bool\_lit>  
 <first\_action\_prop> → **first\_action** : <nonneg\_int\_lit>  
 <type\_idx\_prop> → **type\_idx** : <nonneg\_int\_lit>  
 <next\_prop> → **next** : <nonneg\_int\_lit>  
 | **next** : **none**  
 <lsda\_idx\_prop> → **lsda\_idx** : <nonneg\_int\_lit>  
 <int\_lit> → <digits>  
 | - <digits>  
 <digits> → **/[0-9]+/**  
 <register\_lit> → **r** <digits>  
 <nonneg\_int\_lit> → <digits>

<bool\_lit> → true  
| false

<dw\_pe\_lit> → <dw\_pe\_lit> , <dw\_pe\_lit>  
| DW\_EH\_PE\_ahp  
| DW\_EH\_PE\_alsz  
| DW\_EH\_PE\_ahpalsz  
| DW\_EH\_PE\_ahpalszlsb  
| DW\_EH\_PE\_ahpalszlsb2  
| DW\_EH\_PE\_ahpalszlsb4  
| DW\_EH\_PE\_ahpalszlsb8  
| DW\_EH\_PE\_ahpalszlsb16  
| DW\_EH\_PE\_ahpalszlsb32  
| DW\_EH\_PE\_ahpalszlsb64  
| DW\_EH\_PE\_ahpalszlsb128  
| DW\_EH\_PE\_ahpalszlsb256  
| DW\_EH\_PE\_ahpalszlsb512  
| DW\_EH\_PE\_ahpalszlsb1024  
| DW\_EH\_PE\_ahpalszlsb2048  
| DW\_EH\_PE\_ahpalszlsb4096  
| DW\_EH\_PE\_ahpalszlsb8192  
| DW\_EH\_PE\_ahpalszlsb16384  
| DW\_EH\_PE\_ahpalszlsb32768  
| DW\_EH\_PE\_ahpalszlsb65536  
| DW\_EH\_PE\_ahpalszlsb131072  
| DW\_EH\_PE\_ahpalszlsb262144  
| DW\_EH\_PE\_ahpalszlsb524288  
| DW\_EH\_PE\_ahpalszlsb1048576  
| DW\_EH\_PE\_ahpalszlsb2097152  
| DW\_EH\_PE\_ahpalszlsb4194304  
| DW\_EH\_PE\_ahpalszlsb8388608  
| DW\_EH\_PE\_ahpalszlsb16777216  
| DW\_EH\_PE\_ahpalszlsb33554432  
| DW\_EH\_PE\_ahpalszlsb67108864  
| DW\_EH\_PE\_ahpalszlsb134217728  
| DW\_EH\_PE\_ahpalszlsb268435456  
| DW\_EH\_PE\_ahpalszlsb536870912  
| DW\_EH\_PE\_ahpalszlsb1073741824  
| DW\_EH\_PE\_ahpalszlsb2147483648  
| DW\_EH\_PE\_ahpalszlsb4294967296  
| DW\_EH\_PE\_ahpalszlsb8589934592  
| DW\_EH\_PE\_ahpalszlsb17179869184  
| DW\_EH\_PE\_ahpalszlsb34359738368  
| DW\_EH\_PE\_ahpalszlsb68719476736  
| DW\_EH\_PE\_ahpalszlsb137438953472  
| DW\_EH\_PE\_ahpalszlsb274877906944  
| DW\_EH\_PE\_ahpalszlsb549755813888  
| DW\_EH\_PE\_ahpalszlsb1099511627776  
| DW\_EH\_PE\_ahpalszlsb2199023255552  
| DW\_EH\_PE\_ahpalszlsb4398046511104  
| DW\_EH\_PE\_ahpalszlsb8796093022208  
| DW\_EH\_PE\_ahpalszlsb17592186044416  
| DW\_EH\_PE\_ahpalszlsb35184372088832  
| DW\_EH\_PE\_ahpalszlsb70368744177664  
| DW\_EH\_PE\_ahpalszlsb140737488355328  
| DW\_EH\_PE\_ahpalszlsb281474976710656  
| DW\_EH\_PE\_ahpalszlsb562949953421312  
| DW\_EH\_PE\_ahpalszlsb1125899906842624  
| DW\_EH\_PE\_ahpalszlsb2251799813685248  
| DW\_EH\_PE\_ahpalszlsb4503599627370496  
| DW\_EH\_PE\_ahpalszlsb9007199254740992  
| DW\_EH\_PE\_ahpalszlsb18014398509481984  
| DW\_EH\_PE\_ahpalszlsb36028797018963968  
| DW\_EH\_PE\_ahpalszlsb72057594037927936  
| DW\_EH\_PE\_ahpalszlsb144115188075855872  
| DW\_EH\_PE\_ahpalszlsb288230376151711744  
| DW\_EH\_PE\_ahpalszlsb576460752303423488  
| DW\_EH\_PE\_ahpalszlsb1152921504606846976  
| DW\_EH\_PE\_ahpalszlsb2305843009213693952  
| DW\_EH\_PE\_ahpalszlsb4611686018427387904  
| DW\_EH\_PE\_ahpalszlsb9223372036854775808  
| DW\_EH\_PE\_ahpalszlsb18446744073709551616  
| DW\_EH\_PE\_ahpalszlsb36893488147419103232  
| DW\_EH\_PE\_ahpalszlsb73786976294838206464  
| DW\_EH\_PE\_ahpalszlsb147573952589676412928  
| DW\_EH\_PE\_ahpalszlsb295147905179352825856  
| DW\_EH\_PE\_ahpalszlsb590295810358705651712  
| DW\_EH\_PE\_ahpalszlsb1180591620717411303424  
| DW\_EH\_PE\_ahpalszlsb2361183241434822606848  
| DW\_EH\_PE\_ahpalszlsb4722366482869645213696  
| DW\_EH\_PE\_ahpalszlsb9444732965739290427392  
| DW\_EH\_PE\_ahpalszlsb18889465931478580854784  
| DW\_EH\_PE\_ahpalszlsb37778931862957161709568  
| DW\_EH\_PE\_ahpalszlsb75557863725914323419136  
| DW\_EH\_PE\_ahpalszlsb151115727451828646838272  
| DW\_EH\_PE\_ahpalszlsb302231454903657293676544  
| DW\_EH\_PE\_ahpalszlsb604462909807314587353088  
| DW\_EH\_PE\_ahpalszlsb1208925819614629174706176  
| DW\_EH\_PE\_ahpalszlsb2417851639229258349412352  
| DW\_EH\_PE\_ahpalszlsb4835703278458516698824704  
| DW\_EH\_PE\_ahpalszlsb9671406556917033397649408  
| DW\_EH\_PE\_ahpalszlsb19342813113834066795298816  
| DW\_EH\_PE\_ahpalszlsb38685626227668133590597632  
| DW\_EH\_PE\_ahpalszlsb77371252455336267181195264  
| DW\_EH\_PE\_ahpalszlsb154742504910672534362390528  
| DW\_EH\_PE\_ahpalszlsb309485009821345068724781056  
| DW\_EH\_PE\_ahpalszlsb618970019642690137449562112  
| DW\_EH\_PE\_ahpalszlsb1237940039285380274899244224  
| DW\_EH\_PE\_ahpalszlsb2475880078570760549798488448  
| DW\_EH\_PE\_ahpalszlsb4951760157141521099596976896  
| DW\_EH\_PE\_ahpalszlsb9903520314283042199193953792  
| DW\_EH\_PE\_ahpalszlsb19807040628566084398387907584  
| DW\_EH\_PE\_ahpalszlsb39614081257132168796775815168  
| DW\_EH\_PE\_ahpalszlsb79228162514264337593551630336  
| DW\_EH\_PE\_ahpalszlsb158456325028528675187103260672  
| DW\_EH\_PE\_ahpalszlsb316912650057057350374206521344  
| DW\_EH\_PE\_ahpalszlsb633825300114114700748413042688  
| DW\_EH\_PE\_ahpalszlsb1267650600228229401496826085376  
| DW\_EH\_PE\_ahpalszlsb2535301200456458802993652170752  
| DW\_EH\_PE\_ahpalszlsb5070602400912917605987304341504  
| DW\_EH\_PE\_ahpalszlsb10141204801825835211974608683008  
| DW\_EH\_PE\_ahpalszlsb20282409603651670423949217366016  
| DW\_EH\_PE\_ahpalszlsb40564819207303340847898434732032  
| DW\_EH\_PE\_ahpalszlsb81129638414606681695796869464064  
| DW\_EH\_PE\_ahpalszlsb162259276829213363391597389328128  
| DW\_EH\_PE\_ahpalszlsb324518553658426726783194778656256  
| DW\_EH\_PE\_ahpalszlsb649037107316853453566389557312512  
| DW\_EH\_PE\_ahpalszlsb1298074214633707107132779154625024  
| DW\_EH\_PE\_ahpalszlsb2596148429267414214265558309250048  
| DW\_EH\_PE\_ahpalszlsb519229685853482842853111661850096  
| DW\_EH\_PE\_ahpalszlsb1038459371706965685706223323700192  
| DW\_EH\_PE\_ahpalszlsb2076918743413931371412446647400384  
| DW\_EH\_PE\_ahpalszlsb4153837486827862742824893294800768  
| DW\_EH\_PE\_ahpalszlsb8307674973655725485649786589601536  
| DW\_EH\_PE\_ahpalszlsb16615349947311450971299571779203072  
| DW\_EH\_PE\_ahpalszlsb33230699894622901942599143558406144  
| DW\_EH\_PE\_ahpalszlsb66461399789245803885198287116812288  
| DW\_EH\_PE\_ahpalszlsb132922799578491607770396574233624576  
| DW\_EH\_PE\_ahpalszlsb265845599156983215540793148467249152  
| DW\_EH\_PE\_ahpalszlsb531691198313966431081586296934498304  
| DW\_EH\_PE\_ahpalszlsb1063382396627932862163172593868996608  
| DW\_EH\_PE\_ahpalszlsb2126764793255865724326345187737993216  
| DW\_EH\_PE\_ahpalszlsb4253529586511731448652690375475986432  
| DW\_EH\_PE\_ahpalszlsb8507059173023462897305380750951972864  
| DW\_EH\_PE\_ahpalszlsb17014118346046925794610761501903945728  
| DW\_EH\_PE\_ahpalszlsb34028236692093851589221523003807891456  
| DW\_EH\_PE\_ahpalszlsb68056473384187703178443046007615782912  
| DW\_EH\_PE\_ahpalszlsb136112946768375406356886092015231565824  
| DW\_EH\_PE\_ahpalszlsb27222589353675081271377218403046313152  
| DW\_EH\_PE\_ahpalszlsb54445178707350162542754436806092626304  
| DW\_EH\_PE\_ahpalszlsb108890357414700325085508873612185252608  
| DW\_EH\_PE\_ahpalszlsb217780714829400650171017747224370505216  
| DW\_EH\_PE\_ahpalszlsb435561429658801300342035494448741010432  
| DW\_EH\_PE\_ahpalszlsb871122859317602600684070988897482020864  
| DW\_EH\_PE\_ahpalszlsb1742245718635205201368141977794844041728  
| DW\_EH\_PE\_ahpalszlsb3484491437270410402736283955589688083456  
| DW\_EH\_PE\_ahpalszlsb6968982874540820805472567911179376166912  
| DW\_EH\_PE\_ahpalszlsb13937965749081641610945135822358752333824  
| DW\_EH\_PE\_ahpalszlsb27875931498163283221890271644717504667648  
| DW\_EH\_PE\_ahpalszlsb5575186299632656644378054328943500933536  
| DW\_EH\_PE\_ahpalszlsb11150372599265313288756108657887001867072  
| DW\_EH\_PE\_ahpalszlsb22300745198530626577512217315774003734144  
| DW\_EH\_PE\_ahpalszlsb44601490397061253155024434631548007468288  
| DW\_EH\_PE\_ahpalszlsb89202980794122506310048869263096014936576  
| DW\_EH\_PE\_ahpalszlsb178405961588245012620097784526192029873152  
| DW\_EH\_PE\_ahpalszlsb356811923176490025240195569052384059746304  
| DW\_EH\_PE\_ahpalszlsb713623846352980050480391138104768119492608  
| DW\_EH\_PE\_ahpalszlsb1427247692705960100960782276209536238985216  
| DW\_EH\_PE\_ahpalszlsb2854495385411920201921564552419072477970432  
| DW\_EH\_PE\_ahpalszlsb5708990770823840403843129104838144955840864  
| DW\_EH\_PE\_ahpalszlsb11417981541647680807686258209676289911681728  
| DW\_EH\_PE\_ahpalszlsb22835963083295361615372516419352579823363456  
| DW\_EH\_PE\_ahpalszlsb45671926166590723230745032838705159646726912  
| DW\_EH\_PE\_ahpalszlsb91343852333181446461490065677410319293453824  
| DW\_EH\_PE\_ahpalszlsb18268770466636289292298013135482063858690752  
| DW\_EH\_PE\_ahpalszlsb3653754093327257858459602627096412771381504  
| DW\_EH\_PE\_ahpalszlsb7307508186654515716919205254192825542763008  
| DW\_EH\_PE\_ahpalszlsb14615016373309031433838410508385651085526016  
| DW\_EH\_PE\_ahpalszlsb29230032746618062867676821016771302171052032  
| DW\_EH\_PE\_ahpalszlsb58460065493236125735353642033542604342104064  
| DW\_EH\_PE\_ahpalszlsb116920130986472251470707284067085208684208128  
| DW\_EH\_PE\_ahpalszlsb233840261972944502941414568134170417368416256  
| DW\_EH\_PE\_ahpalszlsb467680523945889005882829136268340834736832512  
| DW\_EH\_PE\_ahpalszlsb935361047891778011765658272536681673473665024  
| DW\_EH\_PE\_ahpalszlsb1870722095783556023531316545073363346947330048  
| DW\_EH\_PE\_ahpalszlsb3741444191567112047062633090146726693894660096  
| DW\_EH\_PE\_ahpalszlsb7482888383134224094125266180293533387789320192  
| DW\_EH\_PE\_ahpalszlsb1496577676626844818825052376058706677557840384  
| DW\_EH\_PE\_ahpalszlsb2993155353253689637650104752117413355115680768  
| DW\_EH\_PE\_ahpalszlsb5986310706507379275300209504234826710231331536  
| DW\_EH\_PE\_ahpalszlsb11972621413014758550600419008469653420462663072  
| DW\_EH\_PE\_ahpalszlsb23945242826029517101200838016939306840925326144  
| DW\_EH\_PE\_ahpalszlsb47890485652059034202401676033878613681850652288  
| DW\_EH\_PE\_ahpalszlsb95780971304118068404803352067757227363701304576  
| DW\_EH\_PE\_ahpalszlsb191561942608236136809606704135514454727402609152  
| DW\_EH\_PE\_ahpalszlsb383123885216472273619213408271028909454805218304  
| DW\_EH\_PE\_ahpalszlsb766247770432944547238426816542057818909610436608  
| DW\_EH\_PE\_ahpalszlsb1532495540865889094476853733084115637819220873216  
| DW\_EH\_PE\_ahpalszlsb3064991081731778188953667466168231275638441746432  
| DW\_EH\_PE\_ahpalszlsb6129982163463556377907334932336462551276883492864  
| DW\_EH\_PE\_ahpalszlsb12259964326927112755814669864672925102553669895728  
| DW\_EH\_PE\_ahpalszlsb24519928653854225511629339729345850205107397791552  
| DW\_EH\_PE\_ahpalszlsb49039857307708451023258679458691700401021559583104  
| DW\_EH\_PE\_ahpalszlsb9807971461541690204651735891738340080204311916608  
| DW\_EH\_PE\_ahpalszlsb19615942922883380409303471823476680160408623833216  
| DW\_EH\_PE\_ahpalszlsb39231885845766760818606943646953360320817246666432  
| DW\_EH\_PE\_ahpalszlsb78463771691533521637213887293906720641634493333264  
| DW\_EH\_PE\_ahpalszlsb156927543383067043274427774587813441283268986666528  
| DW\_EH\_PE\_ahpalszlsb313855086766134086548855549175626882566537733333056  
| DW\_EH\_PE\_ahpalszlsb627710173532268173097711098351253761133107466666112  
| DW\_EH\_PE\_ahpalszlsb1255420347064536346195422196702575226662149333322224  
| DW\_EH\_PE\_ahpalszlsb2510840694129072692390844393405150453324298666644448  
| DW\_EH\_PE\_ahpalszlsb5021681388258145384781688786810300906648597333288896  
| DW\_EH\_PE\_ahpalszlsb10043362776516290769563375736220601813297194666577792  
| DW\_EH\_PE\_ahpalszlsb200867255530325815391267514724412036265943893331555584  
| DW\_EH\_PE\_ahpalszlsb401734511060651630782535029448824072531887786663111104  
| DW\_EH\_PE\_ahpalszlsb80346902212130326156507005889768014506377557332222208  
| DW\_EH\_PE\_ahpalszlsb160693804424260652313014011779536029012755114664444416  
| DW\_EH\_PE\_ahpalszlsb321387608848521304626028023559072058025510229328888832  
| DW\_EH\_PE\_ahpalszlsb642775217697042609252056047118144116051020458657777664  
| DW\_EH\_PE\_ahpalszlsb1285550435394085218504112094236288232102040917315555328  
| DW\_EH\_PE\_ahpalszlsb2571100870788170437008224188472576464204081834631111056  
| DW\_EH\_PE\_ahpalszlsb51422017415763408740164483769451492884081636692622222112  
| DW\_EH\_PE\_ahpalszlsb10284403483152681748032896753890298576816273338444444224  
| DW\_EH\_PE\_ahpalszlsb20568806966305363496065793487780971553632546676888888448  
| DW\_EH\_PE\_ahpalszlsb4113761393261072699213158697556194310726513353377777696  
| DW\_EH\_PE\_ahpalszlsb8227522786522145398426317395112388621442626706755555392  
| DW\_EH\_PE\_ahpalszlsb16455045573044290796852634790224777242885253413511111784  
| DW\_EH\_PE\_ahpalszlsb3291009114608858159370526958044954485761006822702222368  
| DW\_EH\_PE\_ahpalszlsb6582018229217716318741053916089089715522013645404444736  
| DW\_EH\_PE\_ahpalszlsb131640364584354326374821078321781794310440272908088884736  
| DW\_EH\_PE\_ahpalszlsb2632807291687086527496421566435635886208805458161777794736  
| DW\_EH\_PE\_ahpalszlsb5265614583374173054992843132871271772417611111194736  
| DW\_EH\_PE\_ahpalszlsb10531229166748346109985662665742543444232222223894736  
| DW\_EH\_PE\_ahpalszlsb210624583334966922199713253314848868884644444477894736  
| DW\_EH\_PE\_ahpalszlsb42124916666993384439942650662969777777888888957894736  
| DW\_EH\_PE\_ahpalszlsb84249833333986768879885301325939555557777779157894736  
| DW\_EH\_PE\_ahpalszlsb1684996666679735377597706065198791111155555583157894736  
| DW\_EH\_PE\_ahpalszlsb33699933333594707551954121303974222223111111663157894736  
| DW\_EH\_PE\_ahpalszlsb673998666671894151039082426079484444462222223263157894736  
| DW\_EH\_PE\_ahpalszlsb13479973333437890220781648521589688888444444125263157894736  
| DW\_EH\_PE\_ahpalszlsb26959946666875780441563290443779377778888888505263157894736  
| DW\_EH\_PE\_ahpalszlsb53919893333755560883112580887557755577777770105263157894736  
| DW\_EH\_PE\_ahpalszlsb1078397866671111217662251617751151111555555402105263157894736  
| DW\_EH\_PE\_ahpalszlsb2156795733342222435324503235522222311111118042105263157894736  
| DW\_EH\_PE\_ahpalszlsb43135914666844448706490064710444446222222360842105263157894736  
| DW\_EH\_PE\_ahpalszlsb862718293336888974129801294208888844444412016842105263157894736  
| DW\_EH\_PE\_ahpalszlsb1725436586673777948259602584417777888888240336842105263157894736  
| DW\_EH\_PE\_ahpalszlsb34508731733475558965192051688355577777774806736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb690174634669511179303841033767111555555480136736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb138034926933902235860768206753422311111119602734736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb276069853867804471721536413506844222222320054694736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb5521397077356089434430728270136884444446401093894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb110427941547121788688614565402737788888812021877894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb220855883094243577377229130804755777777240437557894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb441711766188487154754458261609511555554808751157894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb8834235323769743095089165232190231111111961542357894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb1766847064753948619017833046380462222223203084736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb35336941295078972380356660927609244444464061694736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb7067388259015794476071332185521888888120123894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb141347765180315889521426643710437777774802477894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb282695530360631779042853287420875555548049557894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb565391060721263558085706574841751111111960991157894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb1130782121442527116171133149763502222223201982357894736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb2261564242885054232342266299527004444446403964736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb4523128485770108464684532599054008888812079294736842105263157894736  
| DW\_EH\_PE\_ahpalszlsb904



| <dw\_cfa\_def\_cfa>  
| <dw\_cfa\_def\_cfa\_sf>  
| <dw\_cfa\_def\_cfa\_register>  
| <dw\_cfa\_def\_cfa\_offset>  
| <dw\_cfa\_def\_cfa\_offset\_sf>  
| <dw\_cfa\_def\_cfa\_expression>  
| <dw\_cfa\_expression>  
| <dw\_cfa\_val\_offset>  
| <dw\_cfa\_val\_offset\_sf>  
| <dw\_cfa\_val\_expression>

<dw\_cfa\_set\_loc> → **DW\_CFA\_set\_loc** <nonneg\_int\_lit>

<dw\_cfa\_advance\_loc> → **DW\_CFA\_advance\_loc** <nonneg\_int\_lit>

<dw\_cfa\_advance\_loc1> → **DW\_CFA\_advance\_loc1** <nonneg\_int\_lit>

<dw\_cfa\_advance\_loc2> → **DW\_CFA\_advance\_loc2** <nonneg\_int\_lit>

<dw\_cfa\_advance\_loc4> → **DW\_CFA\_advance\_loc4** <nonneg\_int\_lit>

<dw\_cfa\_offset> → **DW\_CFA\_offset** <register\_lit><nonneg\_int\_lit>

<dw\_cfa\_offset\_extended> → **DW\_CFA\_offset\_extended** <register\_lit>  
<nonneg\_int\_lit>

<dw\_cfa\_offset\_extended\_sf> → **DW\_CFA\_offset\_extended\_sf**  
<register\_lit> <int\_lit>

<dw\_cfa\_val\_offset> → **DW\_CFA\_val\_offset** <register\_lit> <nonneg\_int\_lit>

<dw\_cfa\_val\_offset\_sf> → **DW\_CFA\_val\_offset\_sf** <register\_lit>  
<int\_lit>

<dw\_cfa\_restore> → **DW\_CFA\_restore** <register\_lit>

<dw\_cfa\_restore\_extended> → **DW\_CFA\_restore\_extended** <register\_lit>

<dw\_cfa\_undefined> → **DW\_CFA\_undefined** <register\_lit>

<dw\_cfa\_same\_value> → **DW\_CFA\_same\_value** <register\_lit>

<dw\_cfa\_register> → **DW\_CFA\_register** <register\_lit> <register\_lit>

<dw\_cfa\_remember\_state> → **DW\_CFA\_remember\_state**

<dw\_cfa\_restore\_state> → **DW\_CFA\_restore\_state**

<dw\_cfa\_def\_cfa> → **DW\_CFA\_def\_cfa** <register\_lit> <nonneg\_int\_lit>

<dw\_cfa\_def\_cfa\_sf> → **DW\_CFA\_def\_cfa\_sf** <register\_lit> <int\_lit>

<dw\_cfa\_def\_cfa\_register> → **DW\_CFA\_def\_cfa\_register** <register\_lit>

<dw\_cfa\_def\_cfa\_offset> → **DW\_CFA\_def\_cfa\_offset** <nonneg\_int\_lit>

<dw\_cfa\_def\_cfa\_offset\_sf> → **DW\_CFA\_def\_cfa\_offset\_sf** <int\_lit>

<dw\_cfa\_def\_cfa\_expression> → **DW\_CFA\_def\_cfa\_expression**  
<expression\_section>

<dw\_cfa\_expression> → **DW\_CFA\_expression** <register\_lit>  
<expression\_section>

<dw\_cfa\_val\_expression> → **DW\_CFA\_val\_expression** <register\_lit>  
<expression\_section>

<expr\_stmt> → <dw\_op\_addr>

| **DW\_OP\_deref**

| <dw\_op\_const1u>

| <dw\_op\_const1s>

| <dw\_op\_const2u>

| <dw\_op\_const2s>

| <dw\_op\_const4u>

| <dw\_op\_const4s>

| <dw\_op\_const8u>

| <dw\_op\_const8s>

| <dw\_op\_constu>

| <dw\_op\_consts>

| **DW\_OP\_dup**

| **DW\_OP\_drop**

| **DW\_OP\_over**

| <dw\_op\_pick>

| **DW\_OP\_swap**

| **DW\_OP\_rot**

| **DW\_OP\_xderef**

| **DW\_OP\_abs**

| **DW\_OP\_and**

| **DW\_OP\_div**

| **DW\_OP\_minus**

| **DW\_OP\_mod**

| **DW\_OP\_mul**

| **DW\_OP\_neg**

| **DW\_OP\_not**

| **DW\_OP\_or**

| **DW\_OP\_plus**

| <dw\_op\_plus\_uconst>

| **DW\_OP\_shl**

| **DW\_OP\_shr**

| **DW\_OP\_shra**

| **DW\_OP\_xor**

| <dw\_op\_skip>

| <dw\_op\_bra>

| **DW\_OP\_eq**

| **DW\_OP\_ge**

| **DW\_OP\_gt**  
 | **DW\_OP\_le**  
 | **DW\_OP\_lt**  
 | **DW\_OP\_ne**  
 | <dw\_op\_litn>  
 | <dw\_op\_regn>  
 | <dw\_op\_bregn>  
 | <dw\_op\_regx>  
 | <dw\_op\_bregx>  
 | <dw\_op\_deref\_size>  
 | <dw\_op\_xderef\_size>  
 | **DW\_OP\_nop**

<label> → <identifier> :

<identifier> → /[a-zA-Z][a-zA-Z0-9]\*/

<dw\_op\_addr> → **DW\_OP\_addr** <nonneg\_int\_lit>

<dw\_op\_const1u> → **DW\_OP\_const1u** <nonneg\_int\_lit>

<dw\_op\_const1s> → **DW\_OP\_const1s** <int\_lit>

<dw\_op\_const2u> → **DW\_OP\_const2u** <nonneg\_int\_lit>

<dw\_op\_const2s> → **DW\_OP\_const2s** <int\_lit>

<dw\_op\_const4u> → **DW\_OP\_const4u** <nonneg\_int\_lit>

<dw\_op\_const4s> → **DW\_OP\_const4s** <int\_lit>

<dw\_op\_const8u> → **DW\_OP\_const8u** <nonneg\_int\_lit>

<dw\_op\_const8s> → **DW\_OP\_const8s** <int\_lit>

<dw\_op\_constu> → **DW\_OP\_constu** <nonneg\_int\_lit>

<dw\_op\_consts> → **DW\_OP\_consts** <int\_lit>

<dw\_op\_pick> → **DW\_OP\_pick** <nonneg\_int\_lit>

<dw\_op\_xderef> → **DW\_OP\_xderef**

<dw\_op\_plus\_uconst> → **DW\_OP\_plus\_uconst**

<dw\_op\_skip> → **DW\_OP\_skip** <int\_lit>

| **DW\_OP\_skip** <identifier>

<dw\_op\_bra> → **DW\_OP\_bra** <int\_lit>

| **DW\_OP\_bra** <identifier>

<dw\_op\_litn> → **DW\_OP\_lit0**

| **DW\_OP\_lit1**

| **DW\_OP\_lit2**

| **DW\_OP\_lit3**

| **DW\_OP\_lit4**

| **DW\_OP\_lit5**

| **DW\_OP\_lit6**

| **DW\_OP\_lit7**

| **DW\_OP\_lit8**

| **DW\_OP\_lit9**

| **DW\_OP\_lit10**

| **DW\_OP\_lit11**

| **DW\_OP\_lit12**

| **DW\_OP\_lit13**

| **DW\_OP\_lit14**

| **DW\_OP\_lit15**

| **DW\_OP\_lit16**

| **DW\_OP\_lit17**

| **DW\_OP\_lit18**

| DW\_OP\_lit19  
| DW\_OP\_lit20  
| DW\_OP\_lit21  
| DW\_OP\_lit22  
| DW\_OP\_lit23  
| DW\_OP\_lit24  
| DW\_OP\_lit25  
| DW\_OP\_lit26  
| DW\_OP\_lit27  
| DW\_OP\_lit28  
| DW\_OP\_lit29  
| DW\_OP\_lit30  
| DW\_OP\_lit31

<dw\_op\_regn> → DW\_OP\_reg0

| DW\_OP\_reg1  
| DW\_OP\_reg2  
| DW\_OP\_reg3  
| DW\_OP\_reg4  
| DW\_OP\_reg5  
| DW\_OP\_reg6  
| DW\_OP\_reg7  
| DW\_OP\_reg8  
| DW\_OP\_reg9  
| DW\_OP\_reg10  
| DW\_OP\_reg11  
| DW\_OP\_reg12  
| DW\_OP\_reg13  
| DW\_OP\_reg14  
| DW\_OP\_reg15  
| DW\_OP\_reg16  
| DW\_OP\_reg17  
| DW\_OP\_reg18  
| DW\_OP\_reg19  
| DW\_OP\_reg20  
| DW\_OP\_reg21  
| DW\_OP\_reg22

| DW\_OP\_reg23  
| DW\_OP\_reg24  
| DW\_OP\_reg25  
| DW\_OP\_reg26  
| DW\_OP\_reg27  
| DW\_OP\_reg28  
| DW\_OP\_reg29  
| DW\_OP\_reg30  
| DW\_OP\_reg31

<dw\_op\_bregn> → **DW\_OP\_breg0** <int\_lit>

| DW\_OP\_breg1 <int\_lit>  
| DW\_OP\_breg2 <int\_lit>  
| DW\_OP\_breg3 <int\_lit>  
| DW\_OP\_breg4 <int\_lit>  
| DW\_OP\_breg5 <int\_lit>  
| DW\_OP\_breg6 <int\_lit>  
| DW\_OP\_breg7 <int\_lit>  
| DW\_OP\_breg8 <int\_lit>  
| DW\_OP\_breg9 <int\_lit>  
| DW\_OP\_breg10 <int\_lit>  
| DW\_OP\_breg11 <int\_lit>  
| DW\_OP\_breg12 <int\_lit>  
| DW\_OP\_breg13 <int\_lit>  
| DW\_OP\_breg14 <int\_lit>  
| DW\_OP\_breg15 <int\_lit>  
| DW\_OP\_breg16 <int\_lit>  
| DW\_OP\_breg17 <int\_lit>  
| DW\_OP\_breg18 <int\_lit>  
| DW\_OP\_breg19 <int\_lit>  
| DW\_OP\_breg20 <int\_lit>  
| DW\_OP\_breg21 <int\_lit>  
| DW\_OP\_breg22 <int\_lit>  
| DW\_OP\_breg23 <int\_lit>  
| DW\_OP\_breg24 <int\_lit>  
| DW\_OP\_breg25 <int\_lit>  
| DW\_OP\_breg26 <int\_lit>

| **DW\_OP\_breg27** <int\_lit>  
| **DW\_OP\_breg28** <int\_lit>  
| **DW\_OP\_breg29** <int\_lit>  
| **DW\_OP\_breg30** <int\_lit>  
| **DW\_OP\_breg31** <int\_lit>

<dw\_op\_regx> → **DW\_OP\_regx** <nonneg\_int\_lit>

<dw\_op\_bregx> → **DW\_OP\_bregx** <nonneg\_int\_lit> <int\_lit>

<dw\_op\_deref\_size> → **DW\_OP\_deref\_size** <nonneg\_int\_lit>

<dw\_op\_xderef\_size> → **DW\_OP\_xderef\_size** <nonneg\_int\_lit>