

System V Application Binary Interface Intel386 Architecture Processor Supplement Version 1.2

Edited by

H.J. Lu¹, David L Kreitzer², Milind Girkar³, Zia Ansari⁴

Based on

System V Application Binary Interface
AMD64 Architecture Processor Supplement

Edited by

H.J. Lu⁵, Michael Matz⁶, Milind Girkar⁷, Jan Hubička⁸, Andreas Jaeger⁹, Mark Mitchell¹⁰

August 7, 2020

¹hongjiu.lu@intel.com

²david.l.kreitzer@intel.com

³milind.girkar@intel.com

⁴zia.ansari@intel.com

⁵hongjiu.lu@intel.com

⁶matz@suse.de

⁷milind.girkar@intel.com

⁸jh@suse.cz

⁹aj@suse.de

¹⁰mark@codesourcery.com

Contents

1	About this Document	7
1.1	Scope	7
1.2	Related Information	8
2	Low Level System Information	9
2.1	Machine Interface	9
2.1.1	Data Representation	9
2.2	Function Calling Sequence	11
2.2.1	Registers	12
2.2.2	The Stack Frame	12
2.2.3	Parameter Passing and Returning Values	13
2.2.4	Variable Argument Lists	18
2.3	Process Initialization	19
2.3.1	Initial Stack and Register State	19
2.3.2	Thread State	22
2.3.3	Auxiliary Vector	22
2.4	DWARF Definition	25
2.4.1	DWARF Release Number	26
2.4.2	DWARF Register Number Mapping	26
2.5	Stack Unwind Algorithm	26
3	Object Files	30
3.1	Sections	30
3.1.1	Special Sections	30
3.1.2	EH_FRAME sections	30
3.2	Symbol Table	35
3.3	Relocation	36
3.3.1	Relocation Types	36

3.4	Program Property	40
4	Libraries	46
4.1	Unwind Library Interface	46
4.1.1	Exception Handler Framework	47
4.1.2	Data Structures	49
4.1.3	Throwing an Exception	52
4.1.4	Exception Object Management	55
4.1.5	Context Management	55
4.1.6	Personality Routine	57
5	Conventions	62
5.1	C++	63
6	Alternate Code Sequences For Security	64
6.1	Code Sequences without PLT	64
6.1.1	Indirect Call via the GOT Slot	64
6.1.2	Thread-Local Storage without PLT	66
7	Intel MPX Extension	70
7.1	Parameter Passing and Returning of Values	70
7.1.1	Bounds Passing	70
7.1.2	Returning of Bounds	71
8	Intel CET Extension	72
8.1	Program Loading	72
8.1.1	Process GNU_PROPERTY_X86_FEATURE_1_IBT	72
8.1.2	Process GNU_PROPERTY_X86_FEATURE_1_SHSTK	73
8.2	Dynamic Linking	73
A	Linker Optimization	75
A.1	Combine GOTPLT and GOT Slots	75
A.2	Optimize R_386_GOT32X Relocation	76

List of Tables

2.1	Scalar Types	10
2.2	Stack Frame with Base Pointer	13
2.3	Register Usage	15
2.4	Return Value Locations for Fundamental Data Types	16
2.5	Parameter Passing Example	17
2.6	Register Allocation for Parameter Passing Example	18
2.7	Stack Layout at the Call	18
2.8	x87 Floating-Point Control Word	19
2.9	MXCSR Status Bits	20
2.10	EFLAGS Bits	20
2.11	Initial Process Stack	21
2.12	auxv_t Type Definition	22
2.13	Auxiliary Vector Types	23
2.14	DWARF Register Number Mapping	27
2.15	Pointer Encoding Specification Byte	28
3.1	Special sections	30
3.2	Common Information Entry (CIE)	32
3.3	CIE Augmentation Section Content	33
3.4	Frame Descriptor Entry (FDE)	34
3.5	FDE Augmentation Section Content	35
3.6	Relocation Types	38
3.7	Program Property Type Ranges	41
3.8	Program Property Types	42
3.9	GNU_PROPERTY_X86_FEATURE_1_AND Bit Flags	43
3.10	Bit Flags For X86 Instruction Sets	44
3.11	Bit Flags For X86 Processor Features	45
6.1	General Dynamic Model Code Sequence	67

6.2	Local Dynamic Model Code Sequence	68
6.3	GD -> IE Code Transition	68
6.4	GD -> LE Code Transition	68
6.5	LD -> LE Code Transition	69
A.1	Call, Jmp and Mov Conversion	76
A.2	Test and Binop Conversion	77

List of Figures

3.1	Relocatable Fields	36
6.1	Function Call without PLT (PIC)	65
6.2	Function Call without PLT (Non-PIC)	65
6.3	Function Address without PLT (PIC)	65
6.4	Function Address without PLT (Non-PIC)	66
6.5	___tls_get_addr Call	67
7.1	Bound Register Usage	70
8.1	The IBT-enabled PLT	74
8.2	The Second IBT-enabled PLT	74
A.1	Procedure Linkage Table Entry Via GOTPLT Slot	75
A.2	Procedure Linkage Table Entry Via GOT Slot	76

Revision History

- 1.2 — 2016-XX-XX** Convert load via GOT slot to load immediate. Clarify R_386_GOT32 and R_386_GOT32X relocations to specify how to compute relocation without base register. Alternate code sequences to call external functions without PLT.
- 1.1 — 2015-12-07** Add AVX-512 support. Add linker optimization to combine GOTPLT and GOT slots. Add R_386_GOT32X relocation and linker optimization. Add FS/GS Base addresses to DWARF register number mapping. Add Intel MPX support.
- 1.0 — 2015-02-03** Reformat table of Returning Values.

0.1 — 2015-01-19 Initial release.

Chapter 1

About this Document

This document is a supplement to the existing Intel386 System V Application Binary Interface (ABI) document available at <http://www.sco.com/developers/devspecs/abi386-4.pdf>, which describes the Linux IA-32 ABI for processors compatible with the Intel386 Architecture.

Intel processors released after the Pentium processors (Pentium 4, Intel Core, and later), have introduced new architecture features, particularly new registers and corresponding instructions to operate on the registers, like the MMX, Intel SSE(1-4), and Intel AVX instruction set extensions. The C/C++ programming languages have evolved to allow programmers to use new data types (for example, `__m64`, `__m128`, and `__m256`). Many compilers (including the Intel compiler and GCC) have supported these data types for some time. Other features in tools (for example, the decimal floating point types, 64-bit integers, exception handling, and so on) have also been developed since the original ABI was written.

This document describes the conventions and constraints on the implementation of these new features for interoperability between various tools.

1.1 Scope

This document describes the conventions on the new C/C++ language types (including alignment and parameter passing conventions), the relocation symbols in the object binary, and the exception handling mechanism for Intel386 architecture. Some of this work has been discussed before <http://groups.google.com/group/ia32-abi> or <http://www.akkadia.org/drepper/tls.pdf>. The C++ object model that is expected to be followed is described in <http://www.akkadia.org/drepper/tls.pdf>.

[//mentoreembedded.github.io/cxx-abi/](http://mentoreembedded.github.io/cxx-abi/). In particular, this document specifies the information that compilers have to generate and the library routines that do the frame unwinding for exception handling.

1.2 Related Information

Links to useful documents:

- System V Application Binary Interface, Intel386™ Architecture Processor Supplement Fourth Edition: <http://www.sco.com/developers/devspecs/abi386-4.pdf>
- System V Application Binary Interface, AMD64 Architecture Processor Supplement, Draft Version 0.99.6: <http://www.x86-64.org/documentation/abi.pdf>
- Discussion of Intel processor extensions: <http://groups.google.com/group/ia32-abi>
- ELF Handling of Thread-Local Storage: <http://www.akkadia.org/drepper/tls.pdf>
- Thread-Local Storage Descriptors for IA32 and AMD64/EM64T: <http://www.fsfla.org/~lxoliva/writeups/TLS/RFC-TLSDESC-x86.txt>
- Itanium C++ ABI, Revised March 20, 2001: <http://mentoreembedded.github.io/cxx-abi/>

Chapter 2

Low Level System Information

This section describes the low-level system information for the Intel386 System V ABI.

2.1 Machine Interface

The Intel386 processor architecture and data representation are covered in this section.

2.1.1 Data Representation

Within this specification, the term *byte* refers to a 8-bit object, the term *twobyte* refers to a 16-bit object, the term *fourbyte* refers to a 32-bit object, the term *eightbyte* refers to a 64-bit object, and the term *sixteenbyte* refers to a 128-bit object.¹

Fundamental Types

Table 2.1 shows the correspondence between ISO C scalar types and the processor scalar types. `__float80`, `__float128`, `__m64`, `__m128`, `__m256` and `__m512` types are optional.

¹The Intel386 ABI uses the term *halfword* for a 16-bit object, the term *word* for a 32-bit object, the term *doubleword* for a 64-bit object. But most IA-32 processor specific documentation define a *word* as a 16-bit object, a *doubleword* as a 32-bit object, a *quadword* as a 64-bit object and a *double quadword* as a 128-bit object.

Table 2.1: Scalar Types

Type	C	sizeof	Alignment (bytes)	Intel386 Architecture
Integral	<code>_Bool</code> [†]	1	1	boolean
	<code>char</code>	1	1	signed byte
	<code>signed char</code>			
	<code>unsigned char</code>	1	1	unsigned byte
	<code>short</code>	2	2	signed twobyte
	<code>signed short</code>			
	<code>unsigned short</code>	2	2	unsigned twobyte
	<code>int</code>	4	4	signed fourbyte
	<code>signed int</code>			
	<code>enum</code> ^{†††}			
	<code>unsigned int</code>	4	4	unsigned fourbyte
	<code>long</code>	4	4	signed fourbyte
<code>signed long</code>				
<code>unsigned long</code>	4	4	unsigned fourbyte	
<code>long long</code>	8	4	signed eightbyte	
<code>signed long long</code>				
<code>unsigned long long</code>	8	4	unsigned eightbyte	
Pointer	<code>any-type *</code> <code>any-type (*)()</code>	4	4	unsigned fourbyte
Floating-point	<code>float</code>	4	4	single (IEEE-754)
	<code>double</code>	8	4	double (IEEE-754)
	<code>long double</code> ^{†††}			
	<code>__float80</code> ^{††}	12	4	80-bit extended (IEEE-754)
	<code>long double</code> ^{†††}			
	<code>__float128</code> ^{††}	16	16	128-bit extended (IEEE-754)
Complex Floating-point	<code>_Complex float</code>	8	4	complex single (IEEE-754)
	<code>_Complex double</code>	16	4	complex double (IEEE-754)
	<code>_Complex long double</code> ^{†††}			
	<code>_Complex __float80</code> ^{††}	24	4	complex 80-bit extended (IEEE-754)
	<code>_Complex long double</code> ^{†††}			
	<code>_Complex __float128</code> ^{††}	32	16	complex 128-bit extended (IEEE-754)
Decimal-floating-point	<code>_Decimal32</code>	4	4	32bit BID (IEEE-754R)
	<code>_Decimal64</code>	8	8	64bit BID (IEEE-754R)
	<code>_Decimal128</code>	16	16	128bit BID (IEEE-754R)
Packed	<code>__m64</code> ^{††}	8	8	MMX and 3DNow!
	<code>__m128</code> ^{††}	16	16	SSE and SSE-2
	<code>__m256</code> ^{††}	32	32	AVX
	<code>__m512</code> ^{††}	64	64	AVX-512

[†] This type is called `bool` in C++.

^{††} These types are optional.

^{†††} C++ and some implementations of C permit enums larger than an int. The underlying type is bumped to an unsigned int.

^{††††} The `long double` type is 64-bit, the same as the `double` type, on the Android™ platform. More information on the Android™ platform is available from <http://www.android.com/>.

The 128-bit floating-point type uses a 15-bit exponent, a 113-bit mantissa (the high order significant bit is implicit) and an exponent bias of 16383.²

The 80-bit floating-point type uses a 15 bit exponent, a 64-bit mantissa with an explicit high order significant bit and an exponent bias of 16383.³

The value of `_Alignof(max_align_t)` is 16.

A null pointer (for all types) has the value zero.

The type `size_t` is defined as `unsigned int`.

Booleans, when stored in a memory object, are stored as single byte objects the value of which is always 0 (`false`) or 1 (`true`). When stored in integer registers (except for passing as arguments), all 4 bytes of the register are significant; any nonzero value is considered `true`.

The Intel386 architecture in general does not require all data accesses to be properly aligned. Misaligned data accesses may be slower than aligned accesses but otherwise behave identically. The only exceptions are that `__float128`, `_Complex __float128`, `_Decimal128`, `__m128`, `__m256` and `__m512` must always be aligned properly.

Structures and Unions

Structures and unions assume the alignment of their most strictly aligned component. Each member is assigned to the lowest available offset with the appropriate alignment. The size of any object is always a multiple of the object's alignment.

Structure and union objects can require padding to meet size and alignment constraints. The contents of any padding is undefined.

2.2 Function Calling Sequence

This section describes the standard function calling sequence, including stack frame layout, register usage, parameter passing and so on.

The standard calling sequence requirements apply only to global functions. Local functions that are not reachable from other compilation units may use different conventions. Nevertheless, it is recommended that all functions use the standard calling sequence when possible.

²Initial implementations of the Intel386 architecture are expected to support operations on the 128-bit floating-point type only via software emulation.

³This type is the x87 double extended precision data type.

2.2.1 Registers

The Intel386 architecture provides 8 general purpose 32-bit registers. In addition the architecture provides 8 SSE registers, each 128 bits wide and 8 x87 floating point registers, each 80 bits wide. Each of the x87 floating point registers may be referred to in *MMX* mode as a 64-bit register. All of these registers are global to all procedures active for a given thread.

Intel AVX (Advanced Vector Extensions) provides 8 256-bit wide AVX registers (`%ymm0` - `%ymm7`). The lower 128-bits of `%ymm0` - `%ymm7` are aliased to the respective 128b-bit SSE registers (`%xmm0` - `%xmm7`). Intel AVX-512 provides 8 512-bit wide SIMD registers (`%zmm0` - `%zmm7`). The lower 128-bits of `%zmm0` - `%zmm7` are aliased to the respective 128b-bit SSE registers (`%xmm0` - `%xmm7`). The lower 256-bits of `%zmm0` - `%zmm7` are aliased to the respective 256-bit AVX registers (`%ymm0` - `%ymm7`). For purposes of parameter passing and function return, `%xmmN`, `%ymmN` and `%zmmN` refer to the same register. Only one of them can be used at the same time. We use vector register to refer to either SSE, AVX or AVX-512 register. In addition, Intel AVX-512 also provides 8 vector mask registers (`%k0` - `%k7`), each 64-bit wide.

The CPU shall be in x87 mode upon entry to a function. Therefore, every function that uses the *MMX* registers is required to issue an `emms` or `femms` instruction after using *MMX* registers, before returning or calling another function.⁴ The direction flag `DF` in the `%EFLAGS` register must be clear (set to “forward” direction) on function entry and return. Other user flags have no specified role in the standard calling sequence and are *not* preserved across calls.

The control bits of the `MXCSR` register are callee-saved (preserved across calls), while the status bits are caller-saved (not preserved). The x87 status word register is caller-saved, whereas the x87 control word is callee-saved.

2.2.2 The Stack Frame

In addition to registers, each function has a frame on the run-time stack. This stack grows downwards from high addresses. Table 2.2 shows the stack organization.

The end of the input argument area shall be aligned on a 16 (32 or 64, if `__m256` or `__m512` is passed on stack) byte boundary.⁵ In other words, the

⁴All x87 registers are caller-saved, so callees that make use of the *MMX* registers may use the faster `femms` instruction.

⁵The maximum aligned boundary is the maximum alignment of all variables passed on stack. In C11, variable of type `typedef struct { _Alignas (512) int i; } var_t;` is

Table 2.2: Stack Frame with Base Pointer

Position	Contents	Frame
4n+8 (%ebp)	memory argument fourbyte <i>n</i>	Previous
	...	
8 (%ebp)	memory argument fourbyte 0	Current
4 (%ebp)	return address	
0 (%ebp)	previous %ebp value	
-4 (%ebp)	unspecified	
	...	
0 (%esp)	variable size	

value (`%esp + 4`) is always a multiple of 16 (32 or 64) when control is transferred to the function entry point. The stack pointer, `%esp`, always points to the end of the latest allocated stack frame. ⁶

2.2.3 Parameter Passing and Returning Values

After the argument values have been computed, they are placed either in registers or pushed on the stack.

Passing Parameters

Most parameters are passed on the stack. Parameters are pushed onto the stack in reverse order - the last argument in the parameter list has the highest address, that is, it is stored farthest away from the stack pointer at the time of the call.

Padding may be needed to increase the size of each parameter to enforce alignment according to the values in Table 2.1. There is an exception for `__m64` and `_Decimal64`, which are treated as having an alignment of four for the purposes of parameter passing. Additional padding may be necessary to ensure that the

aligned to 512 bytes.

⁶The conventional use of `%ebp` as a frame pointer for the stack frame may be avoided by using `%esp` (the stack pointer) to index into the stack frame. This technique saves two instructions in the prologue and epilogue and makes one additional general-purpose register (`%ebp`) available.

bottom of the parameter block (closest to the stack pointer) is at an address which is $0 \bmod 16$, to guarantee proper alignment to the callee.

The exceptions to parameters passed on stack are as follows:

- The first three parameters of type `__m64` are passed in `%mm0`, `%mm1`, and `%mm2`.
- The first three parameters of type `__m128` are passed in `%xmm0`, `%xmm1`, and `%xmm2`.⁷

If parameters of type `__m256` are required to be passed on the stack, the stack pointer must be aligned on a $0 \bmod 32$ byte boundary at the time of the call.

If parameters of type `__m512` are required to be passed on the stack, the stack pointer must be aligned on a $0 \bmod 64$ byte boundary at the time of the call.

Returning Values

Table 2.4 lists the location used to return a value for each fundamental data type. Aggregate types (`structs` and `unions`) are always returned in memory.

Functions that return scalar floating-point values in registers return them on the top of the x87 register stack, that is, `%st0`. It is the responsibility of the calling function to pop this value from the stack regardless of whether or not the value is actually used. Failure to do so results in undefined behavior. An implication of this requirement is that functions returning scalar floating-point values must be properly prototyped. Again, failure to do so results in undefined behavior.

Returning Values in Memory

Some fundamental types and all aggregate types are returned in memory. For functions that return a value in memory, the caller passes a pointer to the memory location where the called function must write the return value. This pointer is passed to called function as an implicit first argument. The memory location must be properly aligned according to the rules in section 2.1.1. In addition to writing the return value to the proper location, the called function is responsible for popping the implicit pointer argument off the stack and storing it in `%eax` prior

⁷The SSE, AVX and AVX-512 registers share resources. Therefore, if the first `__m128` parameter gets assigned to `%xmm0`, the first `__m256`/`__m512` parameter after that is assigned to `%ymm1`/`%zmm1` and not `%ymm0`/`%zmm0`.

Table 2.3: Register Usage

Register	Usage	Preserved across function calls
%eax	scratch register; also used to return integer and pointer values from functions; also stores the address of a returned struct or union	No
%ebx	callee-saved register; also used to hold the GOT pointer when making function calls via the PLT	Yes
%ecx	scratch register	No
%edx	scratch register; also used to return the upper 32bits of some 64bit return types	No
%esp	stack pointer	Yes
%ebp	callee-saved register; optionally used as frame pointer	Yes
%esi	callee-saved register	yes
%edi	callee-saved register	yes
%xmm0, %ymm0	scratch registers; also used to pass and return <code>__m128</code> , <code>__m256</code> parameters	No
%xmm1–%xmm2, %ymm1–%ymm2	scratch registers; also used to pass <code>__m128</code> , <code>__m256</code> parameters	No
%xmm3–%xmm7, %ymm3–%ymm7	scratch registers	No
%mm0	scratch register; also used to pass and return <code>__m64</code> parameter	No
%mm1–%mm2	used to pass <code>__m64</code> parameters	No
%mm3–%mm7	scratch registers	No
%k0–%k7	scratch registers	No
%st0	scratch register; also used to return <code>float</code> , <code>double</code> , <code>long double</code> , <code>__float80</code> values	No
%st1–%st7	scratch registers	No
%gs	Reserved for system (as thread specific data register)	No
mxcsr	SSE2 control and status word	partial
x87 SW	x87 status word	No
x87 CW	x87 control word	Yes

Table 2.4: Return Value Locations for Fundamental Data Types

Type	C	Return Value Location
Integral	<code>_Bool</code> <code>char</code> <code>signed char</code> <code>unsigned char</code>	<code>%al</code> The upper 24 bits of <code>%eax</code> are undefined. The caller must not rely on these being set in a predefined way by the called function.
	<code>short</code> <code>signed short</code> <code>unsigned short</code>	<code>%ax</code> The upper 16 bits of <code>%eax</code> are undefined. The caller must not rely on these being set in a predefined way by the called function.
	<code>int</code> <code>signed int</code> <code>enum</code> <code>unsigned int</code> <code>long</code> <code>signed long</code> <code>unsigned long</code>	<code>%eax</code>
	<code>long long</code> <code>signed long long</code> <code>unsigned long long</code>	<code>%edx:%eax</code> The most significant 32 bits are returned in <code>%edx</code> . The least significant 32 bits are returned in <code>%eax</code> .
Pointer	<code>any-type *</code> <code>any-type (*)()</code>	<code>%eax</code>
Floating-point	<code>float</code>	<code>%st0</code>
	<code>double</code>	<code>%st0</code>
	<code>long double</code>	<code>%st0</code>
	<code>__float80</code>	<code>%st0</code>
	<code>__float128</code>	memory
Complex floating-point	<code>_Complex float</code>	<code>%edx:%eax</code> The real part is returned in <code>%eax</code> . The imaginary part is returned in <code>%edx</code> .
	<code>_Complex double</code>	memory
	<code>_Complex long double</code>	memory
	<code>_Complex __float80</code>	memory
	<code>_Complex __float128</code>	memory
Decimal-floating-point	<code>_Decimal32</code>	<code>%eax</code>
	<code>_Decimal64</code>	<code>%edx:%eax</code> The most significant 32 bits are returned in <code>%edx</code> . The least significant 32 bits are returned in <code>%eax</code> .
	<code>_Decimal128</code>	memory
Packed	<code>__m64</code>	<code>%mm0</code>
	<code>__m128</code>	<code>%xmm0</code>
	<code>__m256</code>	<code>%ymm0</code>
	<code>__m512</code>	<code>%zmm0</code>

to returning. The calling function may choose to reference the return value via `%eax` after the function returns.

As an example of the register passing conventions, consider the declarations and the function call shown in Table 2.5. The corresponding register allocation is given in Table 2.6, the stack frame layout given in Table 2.7 shows the frame before calling the function.

Table 2.5: Parameter Passing Example

```
typedef struct {
    int a, b;
    double d;
} structparm;
structparm s;
int i;
__m128 v, x, y;
__m256 w, z;

extern structparm func (int i, __m128 v,
                       structparm s, __m256 w,
                       __m128 x, __m128 y,
                       __m256 z);

func (i, v, s, w, x, y, z);
```

Table 2.6: Register Allocation for Parameter Passing Example

Parameter	Location before the call
Return value pointer	(%esp)
i	4(%esp)
v	%xmm0
s	8(%esp)
w	%ymm1
x	%xmm2
y	32(%esp)
z	64(%esp)

Table 2.7: Stack Layout at the Call

Contents	Length
z	32 bytes
padding	16 bytes
y	16 bytes
padding	8 bytes
s	16 bytes
i	4 bytes
Return value pointer	4 bytes ← %esp (32-byte aligned)

When a value of type `_Bool` is returned or passed in a register or on the stack, bit 0 contains the truth value and bits 1 to 7 shall be zero⁸.

2.2.4 Variable Argument Lists

Some otherwise portable C programs depend on the argument passing scheme, implicitly assuming that all arguments are passed on the stack, and arguments

⁸Other bits are left unspecified, hence the consumer side of those values can rely on it being 0 or 1 when truncated to 8 bit.

appear in increasing order on the stack. Programs that make these assumptions never have been portable, but they have worked on many implementations. However, they do not work on the Intel386 architecture because some arguments are passed in registers. Portable C programs must use the header file `<stdarg.h>` in order to handle variable argument lists.

When a function taking variable-arguments is called, all parameters are passed on the stack, including `__m64`, `__m128` and `__m256`. This rule applies to both named and unnamed parameters. Because parameters are passed differently depending on whether or not the called function takes a variable argument list, it is necessary for such functions to be properly prototyped. Failure to do so results in undefined behavior.

2.3 Process Initialization

2.3.1 Initial Stack and Register State

Special Registers

The Intel386 architecture defines floating point instructions. At process startup the two floating point units, SSE2 and x87, both have all floating-point exception status flags cleared. The status of the control words is as defined in tables 2.8 and 2.9.

Table 2.8: x87 Floating-Point Control Word

Field	Value	Note
RC	0	Round to nearest
PC	11	Double extended precision
PM	1	Precision masked
UM	1	Underflow masked
OM	1	Overflow masked
ZM	1	Zero divide masked
DM	1	De-normal operand masked
IM	1	Invalid operation masked

Table 2.9: MXCSR Status Bits

Field	Value	Note
FZ	0	Do not flush to zero
RC	0	Round to nearest
PM	1	Precision masked
UM	1	Underflow masked
OM	1	Overflow masked
ZM	1	Zero divide masked
DM	1	De-normal operand masked
IM	1	Invalid operation masked
DAZ	0	De-normals are not zero

The EFLAGS register contains the system flags, such as the direction flag and the carry flag. The low 16 bits (FLAGS portion) of EFLAGS are accessible by application software. The state of them at process initialization is shown in table 2.10.

Table 2.10: EFLAGS Bits

Field	Value	Note
DF	0	Direction forward
CF	0	No carry
PF	0	Even parity
AF	0	No auxiliary carry
ZF	0	No zero result
SF	0	Unsigned result
OF	0	No overflow occurred

Stack State

This section describes the machine state that `exec` (BA_OS) creates for new processes. Various language implementations transform this initial program state to the state required by the language standard.

For example, a C program begins executing at a function named `main` declared as:

```
extern int main ( int argc , char *argv[ ] , char* envp[ ] );
```

where

argc is a non-negative argument count

argv is an array of argument strings, with `argv[argc] == 0`

envp is an array of environment strings, terminated by a null pointer.

When `main()` returns its value is passed to `exit()` and if that has been over-ridden and returns, `_exit()` (which must be immune to user interposition).

The initial state of the process stack, i.e. when `_start` is called is shown in table 2.11.

Table 2.11: Initial Process Stack

Purpose	Start Address	Length
Unspecified	High Addresses	
Information block, including argument strings, environment strings, auxiliary information ...		varies
Unspecified		
Null auxiliary vector entry		1 fourbyte
Auxiliary vector entries ...		2 fourbytes each
0		fourbyte
Environment pointers ...		1 fourbyte each
0	$4+4*\text{argc}+\text{\%esp}$	fourbyte
Argument pointers	$4+\text{\%esp}$	argc fourbytes
Argument count	\%esp	fourbyte
Undefined	Low Addresses	

Argument strings, environment strings, and the auxiliary information appear in no specific order within the information block and they need not be compactly allocated.

Only the registers listed below have specified values at process entry:

- %ebp** The content of this register is unspecified at process initialization time, but the user code should mark the deepest stack frame by setting the frame pointer to zero.
- %esp** The stack pointer holds the address of the byte with lowest address which is part of the stack. It is guaranteed to be 16-byte aligned at process entry.
- %edx** a function pointer that the application should register with `atexit (BA_OS)`.

It is unspecified whether the data and stack segments are initially mapped with execute permissions or not. Applications which need to execute code on the stack or data segments should take proper precautions, e.g., by calling `mprotect ()`.

2.3.2 Thread State

New threads inherit the floating-point state of the parent thread and the state is private to the thread thereafter.

2.3.3 Auxiliary Vector

The auxiliary vector is an array of the following structures (ref. table 2.12), interpreted according to the `a_type` member.

Table 2.12: `auxv_t` Type Definition

```
typedef struct
{
    int a_type;
    union {
        long a_val;
        void *a_ptr;
        void (*a_fnc) ();
    } a_un;
} auxv_t;
```

The Intel386 ABI uses the auxiliary vector types defined in table 2.13.

Table 2.13: Auxiliary Vector Types

Name	Value	a_un
AT_NULL	0	ignored
AT_IGNORE	1	ignored
AT_EXECFD	2	a_val
AT_PHDR	3	a_ptr
AT_PHEMT	4	a_val
AT_PHNUM	5	a_val
AT_PAGESZ	6	a_val
AT_BASE	7	a_ptr
AT_FLAGS	8	a_val
AT_ENTRY	9	a_ptr
AT_NOTELF	10	a_val
AT_UID	11	a_val
AT_EUID	12	a_val
AT_GID	13	a_val
AT_EGID	14	a_val
AT_PLATFORM	15	a_ptr
AT_HWCAP	16	a_val
AT_CLKTCK	17	a_val
AT_SECURE	23	a_val
AT_BASE_PLATFORM	24	a_ptr
AT_RANDOM	25	a_ptr
AT_HWCAP2	26	a_val
AT_EXECFN	31	a_ptr

AT_NULL The auxiliary vector has no fixed length; instead its last entry's `a_type` member has this value.

AT_IGNORE This type indicates the entry has no meaning. The corresponding value of `a_un` is undefined.

AT_EXECFD At process creation the system may pass control to an interpreter program. When this happens, the system places either an entry of type `AT_EXECFD` or one of type `AT_PHDR` in the auxiliary vector. The entry

for type `AT_EXECPD` uses the `a_val` member to contain a file descriptor open to read the application program's object file.

AT_PHDR The system may create the memory image of the application program before passing control to the interpreter program. When this happens, the `a_ptr` member of the `AT_PHDR` entry tells the interpreter where to find the program header table in the memory image.

AT_PHERENT The `a_val` member of this entry holds the size, in bytes, of one entry in the program header table to which the `AT_PHDR` entry points.

AT_PHNUM The `a_val` member of this entry holds the number of entries in the program header table to which the `AT_PHDR` entry points.

AT_PAGESZ If present, this entry's `a_val` member gives the system page size, in bytes.

AT_BASE The `a_ptr` member of this entry holds the base address at which the interpreter program was loaded into memory. See "Program Header" in the System V ABI for more information about the base address.

AT_FLAGS If present, the `a_val` member of this entry holds one-bit flags. Bits with undefined semantics are set to zero.

AT_ENTRY The `a_ptr` member of this entry holds the entry point of the application program to which the interpreter program should transfer control.

AT_NOTELF The `a_val` member of this entry is non-zero if the program is in another format than ELF.

AT_UID The `a_val` member of this entry holds the real user id of the process.

AT_EUID The `a_val` member of this entry holds the effective user id of the process.

AT_GID The `a_val` member of this entry holds the real group id of the process.

AT_EGID The `a_val` member of this entry holds the effective group id of the process.

AT_PLATFORM The `a_ptr` member of this entry points to a string containing the platform name.

AT_HWCAP The `a_val` member of this entry contains an bitmask of CPU features. It mask to the value returned by `CPUID 1.EDX`.

AT_CLKTCK The `a_val` member of this entry contains the frequency at which `times()` increments.

AT_SECURE The `a_val` member of this entry contains one if the program is in secure mode (for example started with `suid`). Otherwise zero.

AT_BASE_PLATFORM The `a_ptr` member of this entry points to a string identifying the base architecture platform (which may be different from the platform).

AT_RANDOM The `a_ptr` member of this entry points to 16 securely generated random bytes.

AT_HWCAP2 The `a_val` member of this entry contains the extended hardware feature mask. Currently it is 0, but may contain additional feature bits in the future.

AT_EXECFN The `a_ptr` member of this entry is a pointer to the file name of the executed program.

2.4 DWARF Definition

This section⁹ defines the Debug With Arbitrary Record Format (DWARF) debugging format for the Intel386 processor family. The Intel386 ABI does not define a debug format. However, all systems that do implement DWARF on Intel386 shall use the following definitions.

DWARF is a specification developed for symbolic, source-level debugging. The debugging information format does not favor the design of any compiler or debugger. For more information on DWARF, see *DWARF Debugging Format Standard*, available at: <http://www.dwarfstd.org/>.

⁹This section is structured in a way similar to the PowerPC psABI

2.4.1 DWARF Release Number

The DWARF definition requires some machine-specific definitions. The register number mapping needs to be specified for the Intel386 registers. In addition, starting with version 3 the DWARF specification requires processor-specific address class codes to be defined.

2.4.2 DWARF Register Number Mapping

Table 2.14¹⁰ outlines the register number mapping for the Intel386 processor family.¹¹

2.5 Stack Unwind Algorithm

The stack frames are not self descriptive and where stack unwinding is desirable (such as for exception handling) additional unwind information needs to be generated. The information is stored in an allocatable section `.eh_frame` whose format is identical to `.debug_frame` defined by the DWARF debug information standard, see *DWARF Debugging Information Format*, with the following extensions:

Position independence In order to avoid load time relocations for position independent code, the FDE CIE offset pointer should be stored relative to the start of CIE table entry. Frames using this extension of the DWARF standard must set the CIE identifier tag to 1.

Outgoing arguments area delta To maintain the size of the temporarily allocated outgoing arguments area present on the end of the stack (when using `push` instructions), operation `GNU_ARGS_SIZE` (0x2e) can be used. This operation takes a single `uleb128` argument specifying the current size. This information is used to adjust the stack frame when jumping into the exception handler of the function after unwinding the stack frame. Additionally the CIE Augmentation shall contain an exact specification of the encoding used. It is recommended to use a PC relative encoding whenever possible and adjust the size according to the code model used.

¹⁰The table defines Return Address to have a register number, even though the address is stored in `0(%esp)` and not in a physical register.

¹¹This document does not define mappings for privileged registers.

Table 2.14: DWARF Register Number Mapping

Register Name	Number	Abbreviation
General Purpose Register EAX	0	%eax
General Purpose Register ECX	1	%ecx
General Purpose Register EDX	2	%edx
General Purpose Register EBX	3	%ebx
Stack Pointer Register ESP	4	%esp
Frame Pointer Register EBP	5	%ebp
General Purpose Register ESI	6	%esi
General Purpose Register EDI	7	%edi
Return Address RA	8	
Flag Register	9	%EFLAGS
Reserved	10	
Floating Point Registers 0–7	11-18	%st0–%st7
Reserved	19-20	
Vector Registers 0–7	21-28	%xmm0–%xmm7
MMX Registers 0–7	29-36	%mm0–%mm7
Media Control and Status	39	%mxcsr
Segment Register ES	40	%es
Segment Register CS	41	%cs
Segment Register SS	42	%ss
Segment Register DS	43	%ds
Segment Register FS	44	%fs
Segment Register GS	45	%gs
Reserved	46-47	
Task Register	48	%tr
LDT Register	49	%ldtr
Reserved	50-92	
FS Base address	93	%fs.base
GS Base address	94	%gs.base

Table 2.15: Pointer Encoding Specification Byte

Mask	Meaning
0x1	Values are stored as <code>uleb128</code> or <code>sleb128</code> type (according to flag 0x8)
0x2	Values are stored as 2 bytes wide integers (<code>udata2</code> or <code>sdata2</code>)
0x3	Values are stored as 4 bytes wide integers (<code>udata4</code> or <code>sdata4</code>)
0x4	Values are stored as 8 bytes wide integers (<code>udata8</code> or <code>sdata8</code>)
0x8	Values are signed
0x10	Values are PC relative
0x20	Values are text section relative
0x30	Values are data section relative
0x40	Values are relative to the start of function

CIE Augmentations: The augmentation field is formatted according to the augmentation field formatting string stored in the CIE header.

The string may contain the following characters:

- z** Indicates that a `uleb128` is present determining the size of the augmentation section.
- L** Indicates the encoding (and thus presence) of an LSDA pointer in the FDE augmentation.
The data field consist of single byte specifying the way pointers are encoded. It is a mask of the values specified by the table 2.15.
The default DWARF pointer encoding (direct 4-byte absolute pointers) is represented by value 0.
- R** Indicates a non-default pointer encoding for FDE code pointers. The formatting is represented by a single byte in the same way as in the 'L' command.
- P** Indicates the presence and an encoding of a language personality routine in the CIE augmentation. The encoding is represented by a single byte in the same way as in the 'L' command followed by a pointer to the personality function encoded by the specified encoding.

When the augmentation is present, the first command must always be 'z' to allow easy skipping of the information.

In order to simplify manipulation of the unwind tables, the runtime library provide higher level API to stack unwinding mechanism, for details see section [4.1](#).

Chapter 3

Object Files

3.1 Sections

3.1.1 Special Sections

Table 3.1: Special sections

Name	Type	Attributes
<code>.eh_frame</code>	<code>SHT_PROGBITS</code>	<code>SHF_ALLOC</code>

.eh_frame This section holds the unwind function table. The contents are described in Section 3.1.2 of this document.

3.1.2 EH_FRAME sections

The call frame information needed for unwinding the stack is output into one section named `.eh_frame`. An `.eh_frame` section consists of one or more subsections. Each subsection contains a CIE (Common Information Entry) followed by varying number of FDEs (Frame Descriptor Entry). A FDE corresponds to an explicit or compiler generated function in a compilation unit, all FDEs can access the CIE that begins their subsection for data. If the code for a function is not one contiguous block, there will be a separate FDE for each contiguous sub-piece.

If an object file contains C++ template instantiations there shall be a separate CIE immediately preceding each FDE corresponding to an instantiation.

Using the preferred encoding specified below, the `.eh_frame` section can be entirely resolved at link time and thus can become part of the text segment.

`EH_PE` encoding below refers to the pointer encoding as specified in the enhanced LSB Chapter 7 for `Eh_Frame_Hdr`.

Table 3.2: Common Information Entry (CIE)

Field	Length (byte)	Description
Length	4	Length of the CIE (not including this 4-byte field)
CIE id	4	Value 0 for <code>.eh_frame</code> (used to distinguish CIEs and FDEs when scanning the section)
Version	1	Value One (1)
CIE Augmentation String	string	Null-terminated string with legal values being "" or 'z' optionally followed by single occurrences of 'P', 'L', or 'R' in any order. The presence of character(s) in the string dictates the content of field 8, the Augmentation Section. Each character has one or two associated operands in the AS (see table 3.3 for which ones). Operand order depends on position in the string ('z' must be first).
Code Align Factor	uleb128	To be multiplied with the "Advance Location" instructions in the Call Frame Instructions
Data Align Factor	sleb128	To be multiplied with all offsets in the Call Frame Instructions
Ret Address Reg	1/uleb128	A "virtual" register representation of the return address. In Dwarf V2, this is a byte, otherwise it is uleb128. It is a byte in gcc 3.3.x
Optional CIE Augmentation Section	varying	Present if Augmentation String in Augmentation Section field 4 is not 0. See table 3.3 for the content.
Optional Call Frame Instructions	varying	

Table 3.3: CIE Augmentation Section Content

Char	Operands	Length (byte)	Description
z	size	uleb128	Length of the remainder of the Augmentation Section
P	personality_enc	1	Encoding specifier - preferred value is a pc-relative, signed 4-byte
	personality routine	(encoded)	Encoded pointer to personality routine (actually to the PLT entry for the personality routine)
R	code_enc	1	Non-default encoding for the code-pointers (FDE members <code>initial_location</code> and <code>address_range</code> and the operand for <code>DW_CFA_set_loc</code>) - preferred value is pc-relative, signed 4-byte
L	lsda_enc	1	FDE augmentation bodies may contain LSDA pointers. If so they are encoded as specified here - preferred value is pc-relative, signed 4-byte possibly indirect thru a GOT entry

Table 3.4: Frame Descriptor Entry (FDE)

Field	Length (byte)	Description
Length	4	Length of the FDE (not including this 4-byte field)
CIE pointer	4	Distance from this field to the nearest preceding CIE (the value is subtracted from the current address). This value can never be zero and thus can be used to distinguish CIE's and FDE's when scanning the <code>.eh_frame</code> section
Initial Location	var	Reference to the function code corresponding to this FDE. If 'R' is missing from the CIE Augmentation String, the field is an 8-byte absolute pointer. Otherwise, the corresponding <code>EH_PE</code> encoding in the CIE Augmentation Section is used to interpret the reference
Address Range	var	Size of the function code corresponding to this FDE. If 'R' is missing from the CIE Augmentation String, the field is an 8-byte unsigned number. Otherwise, the size is determined by the corresponding <code>EH_PE</code> encoding in the CIE Augmentation Section (the value is always absolute)
Optional FDE Augmentation Section	var	Present if CIE Augmentation String is non-empty. See table 3.5 for the content.
Optional Call Frame Instructions	var	

Table 3.5: FDE Augmentation Section Content

Char	Operands	Length (byte)	Description
z	length	uleb128	Length of the remainder of the Augmentation Section
L	LSDA	var	LSDA pointer, encoded in the format specified by the corresponding operand in the CIE's augmentation body. (only present if length > 0).

The existence and size of the optional call frame instruction area must be computed based on the overall size and the offset reached while scanning the preceding fields of the CIE or FDE.

The overall size of a `.eh_frame` section is given in the ELF section header. The only way to determine the number of entries is to scan the section until the end, counting entries as they are encountered.

3.2 Symbol Table

The `STT_GNU_IFUNC`¹ symbol type is optional. It is the same as `STT_FUNC` except that it always points to a function or piece of executable code which takes no arguments and returns a function pointer. If an `STT_GNU_IFUNC` symbol is referred to by a relocation, then evaluation of that relocation is delayed until load-time. The value used in the relocation is the function pointer returned by an invocation of the `STT_GNU_IFUNC` symbol.

The purpose of the `STT_GNU_IFUNC` symbol type is to allow the run-time to select between multiple versions of the implementation of a specific function. The selection made in general will take the currently available hardware into account and select the most appropriate version.

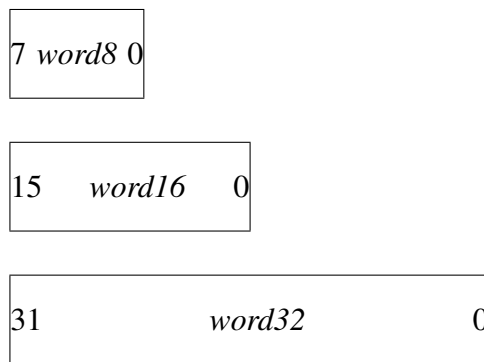
¹It is specified in **Linux Extensions to gABI** at <https://github.com/hjl-tools/linux-abi>

3.3 Relocation

3.3.1 Relocation Types

Figure 3.3.1 shows the allowed relocatable fields.

Figure 3.1: Relocatable Fields



<i>word8</i>	This specifies a 8-bit field occupying 1 byte.
<i>word16</i>	This specifies a 16-bit field occupying 2 bytes with arbitrary byte alignment. These values use the same byte order as other word values in the Intel386 architecture.
<i>word32</i>	This specifies a 32-bit field occupying 4 bytes with arbitrary byte alignment. These values use the same byte order as other word values in the Intel386 architecture.

The following notations are used for specifying relocations in table 3.6:

- A** Represents the addend used to compute the value of the relocatable field.
- B** Represents the base address at which a shared object has been loaded into memory during execution. Generally, a shared object is built with a 0 base virtual address, but the execution address will be different.

- G** Represents the offset into the global offset table at which the relocation entry's symbol will reside during execution.
- GOT** Represents the address of the global offset table.
- L** Represents the place (section offset or address) of the Procedure Linkage Table entry for a symbol.
- P** Represents the place (section offset or address) of the storage unit being relocated (computed using `r_offset`).
- S** Represents the value of the symbol whose index resides in the relocation entry.
- Z** Represents the size of the symbol whose index resides in the relocation entry.

Table 3.6: Relocation Types

Name	Value	Field	Calculation
R_386_NONE	0	none	none
R_386_32	1	word32	S + A
R_386_PC32	2	word32	S + A - P
R_386_GOT32	3	word32	G + A - GOT / G + A [†]
R_386_PLT32	4	word32	L + A - P
R_386_COPY	5	none	none
R_386_GLOB_DAT	6	word32	S
R_386_JUMP_SLOT	7	word32	S
R_386_RELATIVE	8	word32	B + A
R_386_GOTOFF	9	word32	S + A - GOT
R_386_GOTPC	10	word32	GOT + A - P
R_386_TLS_TPOFF	14	word32	
R_386_TLS_IE	15	word32	
R_386_TLS_GOTIE	16	word32	
R_386_TLS_LE	17	word32	
R_386_TLS_GD	18	word32	
R_386_TLS_LDM	19	word32	
R_386_16	20	word16	S + A
R_386_PC16	21	word16	S + A - P
R_386_8	22	word8	S + A
R_386_PC8	23	word8	S + A - P
R_386_TLS_GD_32	24	word32	
R_386_TLS_GD_PUSH	25	word32	
R_386_TLS_GD_CALL	26	word32	
R_386_TLS_GD_POP	27	word32	
R_386_TLS_LDM_32	28	word32	
R_386_TLS_LDM_PUSH	29	word32	
R_386_TLS_LDM_CALL	30	word32	
R_386_TLS_LDM_POP	31	word32	
R_386_TLS_LDO_32	32	word32	
R_386_TLS_IE_32	33	word32	
R_386_TLS_LE_32	34	word32	
R_386_TLS_DTPMOD32	35	word32	
R_386_TLS_DTPOFF32	36	word32	
R_386_TLS_TPOFF32	37	word32	
R_386_SIZE32	38	word32	Z + A
R_386_TLS_GOTDESC	39	word32	
R_386_TLS_DESC_CALL	40	none	none
R_386_TLS_DESC	41	word32	
R_386_IRELATIVE	42	word32	indirect (B + A)
R_386_GOT32X	43	word32	G + A - GOT / G + A [†]

[†] Applied to memory operand without base register when position-independent code is disabled.

R_386_GOT32 and R_386_GOT32X relocations can refer to GOT address, which is a memory operand, or GOT index, which is an immediate operand:

name@GOT It refers to the address of the symbol's global offset table entry.

When it is used without base register and with position-independent code disabled, as in

```
op          name@GOT, %reg
op          %reg, name@GOT
```

it is computed as $G + A$. Otherwise, it is computed as $G + A - GOT$.

For `name@GOT` in:

```
call       *name@GOT(%reg)
jmp        *name@GOT(%reg)
mov        name@GOT(%reg1), %reg2
test      %reg1, name@GOT(%reg2)
binop     name@GOT(%reg1), %reg2
```

as well as

```
call       *name@GOT
jmp        *name@GOT
mov        name@GOT, %reg
test      %reg, name@GOT
binop     name@GOT, %reg
```

where `binop` is one of `adc`, `add`, `and`, `cmp`, `or`, `sbb`, `sub`, `xor` instructions², the `R_386_GOT32X` relocation should be generated, instead of the `R_386_GOT32` relocation. See also section [A.2](#).

`$name@GOT` It refers to the index of the symbol's global offset table entry

```
op          $name@GOT, %reg
```

it is always computed as $G + A - GOT$.

²`mov name@GOT, %eax` must be encoded with opcode `0x8b`, not `0xa0`, to allow linker optimization.

A program or object file using `R_386_8`, `R_386_16`, `R_386_PC16` or `R_386_PC8` relocations is not conformant to this ABI, these relocations are only added for documentation purposes. The `R_386_16`, and `R_386_8` relocations truncate the computed value to 16-bits and 8-bits respectively.

The relocations `R_386_TLS_TPOFF`, `R_386_TLS_IE`, `R_386_TLS_GOTIE`, `R_386_TLS_LE`, `R_386_TLS_GD`, `R_386_TLS_LDM`, `R_386_TLS_GD_32`, `R_386_TLS_GD_PUSH`, `R_386_TLS_GD_CALL`, `R_386_TLS_GD_POP`, `R_386_TLS_LDM_32`, `R_386_TLS_LDM_PUSH`, `R_386_TLS_LDM_CALL`, `R_386_TLS_LDM_POP`, `R_386_TLS_LDO_32`, `R_386_TLS_IE_32`, `R_386_TLS_LE_32`, `R_386_TLS_DTPMOD32`, `R_386_TLS_DTPOFF32` and `R_386_TLS_TPOFF32` are listed for completeness. They are part of the Thread-Local Storage ABI extensions and are documented in the document called “ELF Handling for Thread-Local Storage”³. The relocations `R_386_TLS_GOTDESC`, `R_386_TLS_DESC_CALL` and `R_386_TLS_DESC` are also used for Thread-Local Storage, but are not documented there as of this writing. A description can be found in the document “Thread-Local Storage Descriptors for IA32 and AMD64/EM64T”⁴.

`R_386_IRELATIVE` is similar to `R_386_RELATIVE` except that the value used in this relocation is the program address returned by the function, which takes no arguments, at the address of the result of the corresponding `R_386_RELATIVE` relocation.

One use of the `R_386_IRELATIVE` relocation is to avoid name lookup for the locally defined `STT_GNU_IFUNC` symbols at load-time. Support for this relocation is optional, but is required for the `STT_GNU_IFUNC` symbols.

3.4 Program Property

The following processor-specific program property type ranges⁵ are defined:

³This document is currently available via <http://www.akkadia.org/drepper/tls.pdf>

⁴This document is currently available via <http://www.fsfla.org/~lxoliva/writeups/TLS/RFC-TLSDESC-x86.txt>

⁵See **Linux Extensions to gABI** at <https://github.com/hjl-tools/linux-abi>

Table 3.7: Program Property Type Ranges

Name	Value
GNU_PROPERTY_X86_UINT32_AND_LO	0xc0000002
GNU_PROPERTY_X86_UINT32_AND_HI	0xc0007fff
GNU_PROPERTY_X86_UINT32_OR_LO	0xc0008000
GNU_PROPERTY_X86_UINT32_OR_HI	0xc000ffff
GNU_PROPERTY_X86_UINT32_OR_AND_LO	0xc0010000
GNU_PROPERTY_X86_UINT32_OR_AND_HI	0xc0017fff

The `pr_data` field of each property contains a 4-byte unsigned integer.

GNU_PROPERTY_X86_UINT32_AND_LO..GNU_PROPERTY_X86_UINT32_AND_HI

A bit in the output `pr_data` field is set only if it is set in all relocatable input `pr_data` fields. If all bits in the the output `pr_data` field are zero, this property should be removed from output. If the bit is 1, all input relocatables have the feature. If the bit is 0 or the property is missing, it is unknown whether all input relocatables have the feature.

GNU_PROPERTY_X86_UINT32_OR_LO..GNU_PROPERTY_X86_UINT32_OR_HI

A bit in the output `pr_data` field is set if it is set in any relocatable input `pr_data` fields. If all bits in the the output `pr_data` field are zero, this property should be removed from output. If the bit is 1, some input relocatables have the feature. If the bit is 0 or the property is missing, it is unknown whether any input relocatables have the feature.

GNU_PROPERTY_X86_UINT32_OR_AND_LO..GNU_PROPERTY_X86_UINT32_OR_AND_HI

A bit in the output `pr_data` field is set if it is set in any relocatable input `pr_data` fields and this property is present in all relocatable input files. A missing property implies that its bits have unknown values. When all bits in the the output `pr_data` field are zero, this property should not be removed from output to indicate it has zero in all bits. If the property is in output, all input relocatables have the property. If the bit is 1, some input relocatables have the feature. If the bit is 0, none of input relocatables have the feature.

The following property types are defined:

Table 3.8: Program Property Types

Name	Value
GNU_PROPERTY_X86_FEATURE_1_AND	GNU_PROPERTY_X86_UINT32_AND_LO + 0
GNU_PROPERTY_X86_FEATURE_2_USED	GNU_PROPERTY_X86_UINT32_OR_AND_LO + 1
GNU_PROPERTY_X86_FEATURE_2_NEEDED	GNU_PROPERTY_X86_UINT32_OR_LO + 1
GNU_PROPERTY_X86_ISA_1_USED	GNU_PROPERTY_X86_UINT32_OR_AND_LO + 0
GNU_PROPERTY_X86_ISA_1_NEEDED	GNU_PROPERTY_X86_UINT32_OR_LO + 0

GNU_PROPERTY_X86_FEATURE_1_AND The x86 processor features indicated by the corresponding bits are used in program.

GNU_PROPERTY_X86_FEATURE_2_USED The x86 processor features indicated by the corresponding bits are used in program. Their support in the hardware is optional. Its absence in an x86 ELF binary implies that any x86 processor features may be used. `GNU_PROPERTY_X86_FEATURE_2_USED` can be used to check for features used in the program.

GNU_PROPERTY_X86_FEATURE_2_NEEDED The x86 processor features indicated by the corresponding bits are needed in program and they must be supported by the hardware. Loader may refuse to load the program whose `GNU_PROPERTY_X86_FEATURE_2_NEEDED` features aren't supported by the hardware.

GNU_PROPERTY_X86_ISA_1_NEEDED The x86 instruction sets indicated by the corresponding bits are needed in program and they must be supported by the hardware. Loader may refuse to load the program whose `GNU_PROPERTY_X86_ISA_1_NEEDED` ISAs aren't supported by the hardware.

GNU_PROPERTY_X86_ISA_1_USED The x86 instruction sets indicated by the corresponding bits are used in program. Their support in the hardware is optional. `GNU_PROPERTY_X86_ISA_1_USED` can be used to check for ISAs used in the program.

The following bits are defined for `GNU_PROPERTY_X86_FEATURE_1_AND`:

Table 3.9: GNU_PROPERTY_X86_FEATURE_1_AND Bit Flags

Name	Value
GNU_PROPERTY_X86_FEATURE_1_IBT	1U « 0
GNU_PROPERTY_X86_FEATURE_1_SHSTK	1U « 1

GNU_PROPERTY_X86_FEATURE_1_IBT This indicates that all executable sections are compatible with IBT (see Section 8.1.1) when `endbr64` instruction starts each valid target where an indirect branch instruction can land.

GNU_PROPERTY_X86_FEATURE_1_SHSTK This indicates that all executable sections are compatible with SHSTK (see Section 8.1.2) where return address popped from shadow stack always matches return address popped from normal stack.

Table 3.10: Bit Flags For X86 Instruction Sets

Name	Value
GNU_PROPERTY_X86_ISA_1_CMOV	1U « 0
GNU_PROPERTY_X86_ISA_1_SSE	1U « 1
GNU_PROPERTY_X86_ISA_1_SSE2	1U « 2
GNU_PROPERTY_X86_ISA_1_SSE3	1U « 3
GNU_PROPERTY_X86_ISA_1_SSSE3	1U « 4
GNU_PROPERTY_X86_ISA_1_SSE4_1	1U « 5
GNU_PROPERTY_X86_ISA_1_SSE4_2	1U « 6
GNU_PROPERTY_X86_ISA_1_AVX	1U « 7
GNU_PROPERTY_X86_ISA_1_AVX2	1U « 8
GNU_PROPERTY_X86_ISA_1_FMA	1U « 9
GNU_PROPERTY_X86_ISA_1_AVX512F	1U « 10
GNU_PROPERTY_X86_ISA_1_AVX512CD	1U « 11
GNU_PROPERTY_X86_ISA_1_AVX512ER	1U « 12
GNU_PROPERTY_X86_ISA_1_AVX512PF	1U « 13
GNU_PROPERTY_X86_ISA_1_AVX512VL	1U « 14
GNU_PROPERTY_X86_ISA_1_AVX512DQ	1U « 15
GNU_PROPERTY_X86_ISA_1_AVX512BW	1U « 16
GNU_PROPERTY_X86_ISA_1_AVX512_4FMAPS	1U « 17
GNU_PROPERTY_X86_ISA_1_AVX512_4VNNIW	1U « 18
GNU_PROPERTY_X86_ISA_1_AVX512_BITALG	1U « 19
GNU_PROPERTY_X86_ISA_1_AVX512_IFMA	1U « 20
GNU_PROPERTY_X86_ISA_1_AVX512_VBMI	1U « 21
GNU_PROPERTY_X86_ISA_1_AVX512_VBMI2	1U « 22
GNU_PROPERTY_X86_ISA_1_AVX512_VNNI	1U « 23

Table 3.11: Bit Flags For X86 Processor Features

Name	Value
GNU_PROPERTY_X86_FEATURE_2_X86 [†]	1U « 0
GNU_PROPERTY_X86_FEATURE_2_X87	1U « 1
GNU_PROPERTY_X86_FEATURE_2_MMX	1U « 2
GNU_PROPERTY_X86_FEATURE_2_XMM	1U « 3
GNU_PROPERTY_X86_FEATURE_2_YMM	1U « 4
GNU_PROPERTY_X86_FEATURE_2_ZMM	1U « 5
GNU_PROPERTY_X86_FEATURE_2_FXSR	1U « 6
GNU_PROPERTY_X86_FEATURE_2_XSAVE	1U « 7
GNU_PROPERTY_X86_FEATURE_2_XSAVEOPT	1U « 8
GNU_PROPERTY_X86_FEATURE_2_XSAVEC	1U « 9

[†] This bit should always be set when GNU_PROPERTY_X86_FEATURE_2_USED is used.

Chapter 4

Libraries

4.1 Unwind Library Interface

This section defines the Unwind Library interface¹, expected to be provided by any Intel386 psABI-compliant system. This is the interface on which the C++ ABI exception-handling facilities are built. We assume as a basis the Call Frame Information tables described in the DWARF Debugging Information Format document.

This section is meant to specify a language-independent interface that can be used to provide higher level exception-handling facilities such as those defined by C++.

The unwind library interface consists of at least the following routines:

```
_Unwind_RaiseException ,  
_Unwind_Resume ,  
_Unwind_DeleteException ,  
_Unwind_GetGR ,  
_Unwind_SetGR ,  
_Unwind_GetIP ,  
_Unwind_SetIP ,  
_Unwind_GetRegionStart ,  
_Unwind_GetLanguageSpecificData ,  
_Unwind_ForcedUnwind ,  
_Unwind_GetCFA
```

¹The overall structure and the external interface is derived from the IA-64 UNIX System V ABI

In addition, two data types are defined (`_Unwind_Context` and `_Unwind_Exception`) to interface a calling runtime (such as the C++ runtime) and the above routine. All routines and interfaces behave as if defined `extern "C"`. In particular, the names are not mangled. All names defined as part of this interface have a `"_Unwind_"` prefix.

Lastly, a language and vendor specific personality routine will be stored by the compiler in the unwind descriptor for the stack frames requiring exception processing. The personality routine is called by the unwinder to handle language-specific tasks such as identifying the frame handling a particular exception.

4.1.1 Exception Handler Framework

Reasons for Unwinding

There are two major reasons for unwinding the stack:

- exceptions, as defined by languages that support them (such as C++)
- “forced” unwinding (such as caused by `longjmp` or thread termination)

The interface described here tries to keep both similar. There is a major difference, however.

- In the case where an exception is thrown, the stack is unwound while the exception propagates, but it is expected that the personality routine for each stack frame knows whether it wants to catch the exception or pass it through. This choice is thus delegated to the personality routine, which is expected to act properly for any type of exception, whether “native” or “foreign”. Some guidelines for “acting properly” are given below.
- During “forced unwinding”, on the other hand, an external agent is driving the unwinding. For instance, this can be the `longjmp` routine. This external agent, not each personality routine, knows when to stop unwinding. The fact that a personality routine is not given a choice about whether unwinding will proceed is indicated by the `_UA_FORCE_UNWIND` flag.

To accommodate these differences, two different routines are proposed. `_Unwind_RaiseException` performs exception-style unwinding, under control of the personality routines. `_Unwind_ForcedUnwind`, on the other hand, performs unwinding, but gives an external agent the opportunity to intercept

calls to the personality routine. This is done using a proxy personality routine, that intercepts calls to the personality routine, letting the external agent override the defaults of the stack frame's personality routine.

As a consequence, it is not necessary for each personality routine to know about any of the possible external agents that may cause an unwind. For instance, the C++ personality routine need deal only with C++ exceptions (and possibly disguising foreign exceptions), but it does not need to know anything specific about unwinding done on behalf of `longjmp` or `pthread`s cancellation.

The Unwind Process

The standard ABI exception handling/unwind process begins with the raising of an exception, in one of the forms mentioned above. This call specifies an exception object and an exception class.

The runtime framework then starts a two-phase process:

- In the *search* phase, the framework repeatedly calls the personality routine, with the `_UA_SEARCH_PHASE` flag as described below, first for the current `%eip` and register state, and then unwinding a frame to a new `%eip` at each step, until the personality routine reports either success (a handler found in the queried frame) or failure (no handler) in all frames. It does not actually restore the unwound state, and the personality routine must access the state through the API.
- If the search phase reports a failure, e.g. because no handler was found, it will call `terminate()` rather than commence phase 2.

If the search phase reports success, the framework restarts in the *cleanup* phase. Again, it repeatedly calls the personality routine, with the `_UA_CLEANUP_PHASE` flag as described below, first for the current `%eip` and register state, and then unwinding a frame to a new `%eip` at each step, until it gets to the frame with an identified handler. At that point, it restores the register state, and control is transferred to the user landing pad code.

Each of these two phases uses both the unwind library and the personality routines, since the validity of a given handler and the mechanism for transferring control to it are language-dependent, but the method of locating and restoring previous stack frames is language-independent.

A two-phase exception-handling model is not strictly necessary to implement C++ language semantics, but it does provide some benefits. For example, the first phase allows an exception-handling mechanism to *dismiss* an exception before stack unwinding begins, which allows *presumptive* exception handling (correcting the exceptional condition and resuming execution at the point where it was raised). While C++ does not support presumptive exception handling, other languages do, and the two-phase model allows C++ to coexist with those languages on the stack.

Note that even with a two-phase model, we may execute each of the two phases more than once for a single exception, as if the exception was being thrown more than once. For instance, since it is not possible to determine if a given catch clause will re-throw or not without executing it, the exception propagation effectively stops at each catch clause, and if it needs to restart, restarts at phase 1. This process is not needed for destructors (cleanup code), so the phase 1 can safely process all destructor-only frames at once and stop at the next enclosing catch clause.

For example, if the first two frames unwound contain only cleanup code, and the third frame contains a C++ catch clause, the personality routine in phase 1, does not indicate that it found a handler for the first two frames. It must do so for the third frame, because it is unknown how the exception will propagate out of this third frame, e.g. by re-throwing the exception or throwing a new one in C++.

The API specified by the Intel386 psABI for implementing this framework is described in the following sections.

4.1.2 Data Structures

Reason Codes

The unwind interface uses reason codes in several contexts to identify the reasons for failures or other actions, defined as follows:

```

typedef enum {
    _URC_NO_REASON = 0,
    _URC_FOREIGN_EXCEPTION_CAUGHT = 1,
    _URC_FATAL_PHASE2_ERROR = 2,
    _URC_FATAL_PHASE1_ERROR = 3,
    _URC_NORMAL_STOP = 4,
    _URC_END_OF_STACK = 5,
    _URC_HANDLER_FOUND = 6,
    _URC_INSTALL_CONTEXT = 7,
    _URC_CONTINUE_UNWIND = 8
} _Unwind_Reason_Code;

```

The interpretations of these codes are described below.

Exception Header

The unwind interface uses a pointer to an exception header object as its representation of an exception being thrown. In general, the full representation of an exception object is language- and implementation-specific, but is prefixed by a header understood by the unwind interface, defined as follows:

```

typedef void (*_Unwind_Exception_Cleanup_Fn)
    (_Unwind_Reason_Code reason,
     struct _Unwind_Exception *exc);
struct _Unwind_Exception {
    uint64_t          exception_class;
    _Unwind_Exception_Cleanup_Fn exception_cleanup;
    uint32_t          private_1;
    uint32_t          private_2;
};

```

An `_Unwind_Exception` object must be eightbyte aligned. The first two fields are set by user code prior to raising the exception, and the latter two should never be touched except by the runtime.

The `exception_class` field is a language- and implementation-specific identifier of the kind of exception. It allows a personality routine to distinguish between native and foreign exceptions, for example. By convention, the high 4 bytes indicate the vendor (for instance GNUC), and the low 4 bytes indicate the language. For the C++ ABI described in this document, the low four bytes are C++\0.

The `exception_cleanup` routine is called whenever an exception object needs to be destroyed by a different runtime than the runtime which created the exception object, for instance if a Java exception is caught by a C++ catch handler. In such a case, a reason code (see above) indicates why the exception object needs to be deleted:

`_URC_FOREIGN_EXCEPTION_CAUGHT = 1` This indicates that a different runtime caught this exception. Nested foreign exceptions, or re-throwing a foreign exception, result in undefined behavior.

`_URC_FATAL_PHASE1_ERROR = 3` The personality routine encountered an error during phase 1, other than the specific error codes defined.

`_URC_FATAL_PHASE2_ERROR = 2` The personality routine encountered an error during phase 2, for instance a stack corruption.

Normally, all errors should be reported during phase 1 by returning from `_Unwind_RaiseException`. However, landing pad code could cause stack corruption between phase 1 and phase 2. For a C++ exception, the runtime should call `terminate()` in that case.

The private unwinder state (`private_1` and `private_2`) in an exception object should be neither read by nor written to by personality routines or other parts of the language-specific runtime. It is used by the specific implementation of the unwinder on the host to store internal information, for instance to remember the final handler frame between unwinding phases.

In addition to the above information, a typical runtime such as the C++ runtime will add language-specific information used to process the exception. This is expected to be a contiguous area of memory after the `_Unwind_Exception` object, but this is not required as long as the matching personality routines know how to deal with it, and the `exception_cleanup` routine de-allocates it properly.

Unwind Context

The `_Unwind_Context` type is an opaque type used to refer to a system-specific data structure used by the system unwinder. This context is created and destroyed by the system, and passed to the personality routine during unwinding.

```
struct _Unwind_Context
```

4.1.3 Throwing an Exception

`_Unwind_RaiseException`

```
_Unwind_Reason_Code _Unwind_RaiseException  
( struct _Unwind_Exception *exception_object );
```

Raise an exception, passing along the given exception object, which should have its `exception_class` and `exception_cleanup` fields set. The exception object has been allocated by the language-specific runtime, and has a language-specific format, except that it must contain an `_Unwind_Exception` struct (see Exception Header above). `_Unwind_RaiseException` does not return, unless an error condition is found (such as no handler for the exception, bad stack format, etc.). In such a case, an `_Unwind_Reason_Code` value is returned.

Possibilities are:

`_URC_END_OF_STACK` The unwinder encountered the end of the stack during phase 1, without finding a handler. The unwind runtime will not have modified the stack. The C++ runtime will normally call `uncaught_exception()` in this case.

`_URC_FATAL_PHASE1_ERROR` The unwinder encountered an unexpected error during phase 1, e.g. stack corruption. The unwind runtime will not have modified the stack. The C++ runtime will normally call `terminate()` in this case.

If the unwinder encounters an unexpected error during phase 2, it should return `_URC_FATAL_PHASE2_ERROR` to its caller. In C++, this will usually be `__cxa_throw`, which will call `terminate()`.

The unwind runtime will likely have modified the stack (e.g. popped frames from it) or register context, or landing pad code may have corrupted them. As a result, the the caller of `_Unwind_RaiseException` can make no assumptions about the state of its stack or registers.

_Unwind_ForcedUnwind

```
typedef _Unwind_Reason_Code (*_Unwind_Stop_Fn)
(int version,
 _Unwind_Action actions,
 uint64 exceptionClass,
 struct _Unwind_Exception *exceptionObject,
 struct _Unwind_Context *context,
 void *stop_parameter );
_Unwind_Reason_Code _Unwind_ForcedUnwind
( struct _Unwind_Exception *exception_object,
  _Unwind_Stop_Fn stop,
  void *stop_parameter );
```

Raise an exception for forced unwinding, passing along the given exception object, which should have its `exception_class` and `exception_cleanup` fields set. The exception object has been allocated by the language-specific runtime, and has a language-specific format, except that it must contain an `_Unwind_Exception` struct (see Exception Header above).

Forced unwinding is a single-phase process (phase 2 of the normal exception-handling process). The `stop` and `stop_parameter` parameters control the termination of the unwind process, instead of the usual personality routine query. The `stop` function parameter is called for each unwind frame, with the parameters described for the usual personality routine below, plus an additional `stop_parameter`.

When the `stop` function identifies the destination frame, it transfers control (according to its own, unspecified, conventions) to the user code as appropriate without returning, normally after calling `_Unwind_DeleteException`. If not, it should return an `_Unwind_Reason_Code` value as follows:

_URC_NO_REASON This is not the destination frame. The unwind runtime will call the frame's personality routine with the `_UA_FORCE_UNWIND` and `_UA_CLEANUP_PHASE` flags set in `actions`, and then unwind to the next frame and call the stop function again.

_URC_END_OF_STACK In order to allow `_Unwind_ForcedUnwind` to perform special processing when it reaches the end of the stack, the unwind runtime will call it after the last frame is rejected, with a `NULL` stack pointer

in the context, and the stop function must catch this condition (i.e. by noticing the NULL stack pointer). It may return this reason code if it cannot handle end-of-stack.

`_URC_FATAL_PHASE2_ERROR` The stop function may return this code for other fatal conditions, e.g. stack corruption.

If the stop function returns any reason code other than `_URC_NO_REASON`, the stack state is indeterminate from the point of view of the caller of `_Unwind_ForcedUnwind`. Rather than attempt to return, therefore, the unwind library should return `_URC_FATAL_PHASE2_ERROR` to its caller.

Example: `longjmp_unwind()`

The expected implementation of `longjmp_unwind()` is as follows. The `setjmp()` routine will have saved the state to be restored in its customary place, including the frame pointer. The `longjmp_unwind()` routine will call `_Unwind_ForcedUnwind` with a stop function that compares the frame pointer in the context record with the saved frame pointer. If equal, it will restore the `setjmp()` state as customary, and otherwise it will return `_URC_NO_REASON` or `_URC_END_OF_STACK`.

If a future requirement for two-phase forced unwinding were identified, an alternate routine could be defined to request it, and an actions parameter flag defined to support it.

`_Unwind_Resume`

```
void _Unwind_Resume
(struct _Unwind_Exception *exception_object);
```

Resume propagation of an existing exception e.g. after executing cleanup code in a partially unwound stack. A call to this routine is inserted at the end of a landing pad that performed cleanup, but did not resume normal execution. It causes unwinding to proceed further.

`_Unwind_Resume` should not be used to implement re-throwing. To the unwinding runtime, the catch code that re-throws was a handler, and the previous unwinding session was terminated before entering it. Re-throwing is implemented by calling `_Unwind_RaiseException` again with the same exception object.

This is the only routine in the unwind library which is expected to be called directly by generated code: it will be called at the end of a landing pad in a "landing-pad" model.

4.1.4 Exception Object Management

`_Unwind_DeleteException`

```
void _Unwind_DeleteException
    (struct _Unwind_Exception *exception_object);
```

Deletes the given exception object. If a given runtime resumes normal execution after catching a foreign exception, it will not know how to delete that exception. Such an exception will be deleted by calling `_Unwind_DeleteException`. This is a convenience function that calls the function pointed to by the `exception_cleanup` field of the exception header.

4.1.5 Context Management

These functions are used for communicating information about the unwind context (i.e. the unwind descriptors and the user register state) between the unwind library and the personality routine and landing pad. They include routines to read or set the context record images of registers in the stack frame corresponding to a given unwind context, and to identify the location of the current unwind descriptors and unwind frame.

`_Unwind_GetGR`

```
uint32 _Unwind_GetGR
    (struct _Unwind_Context *context, int index);
```

This function returns the 32-bit value of the given general register. The register is identified by its index as given in table 2.14.

During the two phases of unwinding, no registers have a guaranteed value.

`_Unwind_SetGR`

```
void _Unwind_SetGR
    (struct _Unwind_Context *context,
     int index,
     uint32 new_value);
```

This function sets the 32-bit value of the given register, identified by its index as for `_Unwind_GetGR`.

The behavior is guaranteed only if the function is called during phase 2 of unwinding, and applied to an unwind context representing a handler frame, for

which the personality routine will return `_URC_INSTALL_CONTEXT`. In that case, only registers `%eax` and `%edx` should be used. These scratch registers are reserved for passing arguments between the personality routine and the landing pads.

`_Unwind_GetIP`

```
uint32 _Unwind_GetIP
(struct _Unwind_Context *context);
```

This function returns the 32-bit value of the instruction pointer (IP).

During unwinding, the value is guaranteed to be the address of the instruction immediately following the call site in the function identified by the unwind context. This value may be outside of the procedure fragment for a function call that is known to not return (such as `_Unwind_Resume`).

`_Unwind_SetIP`

```
void _Unwind_SetIP
(struct _Unwind_Context *context,
uint32 new_value);
```

This function sets the value of the instruction pointer (IP) for the routine identified by the unwind context.

The behavior is guaranteed only when this function is called for an unwind context representing a handler frame, for which the personality routine will return `_URC_INSTALL_CONTEXT`. In this case, control will be transferred to the given address, which should be the address of a landing pad.

`_Unwind_GetLanguageSpecificData`

```
uint32 _Unwind_GetLanguageSpecificData
(struct _Unwind_Context *context);
```

This routine returns the address of the language-specific data area for the current stack frame.

This routine is not strictly required: it could be accessed through `_Unwind_GetIP` using the documented format of the DWARF Call Frame Information Tables, but since this work has been done for finding the personality routine in the first place, it makes sense to cache the result in the context. We could also pass it as an argument to the personality routine.

`_Unwind_GetRegionStart`

```
uint32 _Unwind_GetRegionStart  
    (struct _Unwind_Context *context);
```

This routine returns the address of the beginning of the procedure or code fragment described by the current unwind descriptor block.

This information is required to access any data stored relative to the beginning of the procedure fragment. For instance, a call site table might be stored relative to the beginning of the procedure fragment that contains the calls. During unwinding, the function returns the start of the procedure fragment containing the call site in the current stack frame.

`_Unwind_GetCFA`

```
uint32 _Unwind_GetCFA  
    (struct _Unwind_Context *context);
```

This function returns the 32-bit Canonical Frame Address which is defined as the value of `%esp` at the call site in the previous frame. This value is guaranteed to be correct any time the context has been passed to a personality routine or a stop function.

4.1.6 Personality Routine

```
_Unwind_Reason_Code (*__personality_routine)  
    (int version,  
     _Unwind_Action actions,  
     uint64 exceptionClass,  
     struct _Unwind_Exception *exceptionObject,  
     struct _Unwind_Context *context);
```

The personality routine is the function in the C++ (or other language) runtime library which serves as an interface between the system unwind library and language-specific exception handling semantics. It is specific to the code fragment described by an unwind info block, and it is always referenced via the pointer in the unwind info block, and hence it has no psABI-specified name.

Parameters

The personality routine parameters are as follows:

version Version number of the unwinding runtime, used to detect a mis-match between the unwinder conventions and the personality routine, or to provide backward compatibility. For the conventions described in this document, version will be 1.

actions Indicates what processing the personality routine is expected to perform, as a bit mask. The possible actions are described below.

exceptionClass An 8-byte identifier specifying the type of the thrown exception. By convention, the high 4 bytes indicate the vendor (for instance GNUC), and the low 4 bytes indicate the language. For the C++ ABI described in this document, the low four bytes are C++\0. This is not a null-terminated string. Some implementations may use no null bytes.

exceptionObject The pointer to a memory location recording the necessary information for processing the exception according to the semantics of a given language (see the Exception Header section above).

context Unwinder state information for use by the personality routine. This is an opaque handle used by the personality routine in particular to access the frame's registers (see the Unwind Context section above).

return value The return value from the personality routine indicates how further unwind should happen, as well as possible error conditions. See the following section.

Personality Routine Actions

The actions argument to the personality routine is a bitwise OR of one or more of the following constants:

```
typedef int _Unwind_Action;
const _Unwind_Action _UA_SEARCH_PHASE = 1;
const _Unwind_Action _UA_CLEANUP_PHASE = 2;
const _Unwind_Action _UA_HANDLER_FRAME = 4;
const _Unwind_Action _UA_FORCE_UNWIND = 8;
```

UA_SEARCH_PHASE Indicates that the personality routine should check if the current frame contains a handler, and if so return `_URC_HANDLER_FOUND`,

or otherwise return `_URC_CONTINUE_UNWIND`. `_UA_SEARCH_PHASE` cannot be set at the same time as `_UA_CLEANUP_PHASE`.

`_UA_CLEANUP_PHASE` Indicates that the personality routine should perform cleanup for the current frame. The personality routine can perform this cleanup itself, by calling nested procedures, and return `_URC_CONTINUE_UNWIND`. Alternatively, it can setup the registers (including the IP) for transferring control to a "landing pad", and return `_URC_INSTALL_CONTEXT`.

`_UA_HANDLER_FRAME` During phase 2, indicates to the personality routine that the current frame is the one which was flagged as the handler frame during phase 1. The personality routine is not allowed to change its mind between phase 1 and phase 2, i.e. it must handle the exception in this frame in phase 2.

`_UA_FORCE_UNWIND` During phase 2, indicates that no language is allowed to "catch" the exception. This flag is set while unwinding the stack for `long jmp` or during thread cancellation. User-defined code in a catch clause may still be executed, but the catch clause must resume unwinding with a call to `_Unwind_Resume` when finished.

Transferring Control to a Landing Pad

If the personality routine determines that it should transfer control to a landing pad (in phase 2), it may set up registers (including IP) with suitable values for entering the landing pad (e.g. with landing pad parameters), by calling the context management routines above. It then returns `_URC_INSTALL_CONTEXT`.

Prior to executing code in the landing pad, the unwind library restores registers not altered by the personality routine, using the context record, to their state in that frame before the call that threw the exception, as follows. All registers specified as callee-saved by the base ABI are restored, as well as scratch registers `%eax` and `%edx` (see below). Except for those exceptions, scratch (or caller-saved) registers are not preserved, and their contents are undefined on transfer.

The landing pad can either resume normal execution (as, for instance, at the end of a C++ catch), or resume unwinding by calling `_Unwind_Resume` and passing it the `exceptionObject` argument received by the personality routine. `_Unwind_Resume` will never return.

`_Unwind_Resume` should be called if and only if the personality routine did not return `_Unwind_HANDLER_FOUND` during phase 1. As a result, the unwinder can allocate resources (for instance memory) and keep track of them in the exception object reserved words. It should then free these resources before transferring control to the last (handler) landing pad. It does not need to free the resources before entering non-handler landing-pads, since `_Unwind_Resume` will ultimately be called.

The landing pad may receive arguments from the runtime, typically passed in registers set using `_Unwind_SetGR` by the personality routine. For a landing pad that can call to `_Unwind_Resume`, one argument must be the `exceptionObject` pointer, which must be preserved to be passed to `_Unwind_Resume`.

The landing pad may receive other arguments, for instance a switch value indicating the type of the exception. Two scratch registers are reserved for this use (`%eax` and `%edx`).

Rules for Correct Inter-Language Operation

The following rules must be observed for correct operation between languages and/or run times from different vendors:

An exception which has an unknown class must not be altered by the personality routine. The semantics of foreign exception processing depend on the language of the stack frame being unwound. This covers in particular how exceptions from a foreign language are mapped to the native language in that frame.

If a runtime resumes normal execution, and the caught exception was created by another runtime, it should call `_Unwind_DeleteException`. This is true even if it understands the exception object format (such as would be the case between different C++ run times).

A runtime is not allowed to catch an exception if the `_UA_FORCE_UNWIND` flag was passed to the personality routine.

Example: Foreign Exceptions in C++. In C++, foreign exceptions can be caught by a `catch(...)` statement. They can also be caught as if they were of a `__foreign_exception` class, defined in `<exception>`. The `__foreign_exception` may have subclasses, such as `__java_exception` and `__ada_exception`, if the runtime is capable of identifying some of the foreign languages.

The behavior is undefined in the following cases:

- A `__foreign_exception` catch argument is accessed in any way (including taking its address).
- A `__foreign_exception` is active at the same time as another exception (either there is a nested exception while catching the foreign exception, or the foreign exception was itself nested).
- `uncaught_exception()`, `set_terminate()`, `set_unexpected()`, `terminate()`, or `unexpected()` is called at a time a foreign exception exists (for example, calling `set_terminate()` during unwinding of a foreign exception).

All these cases might involve accessing C++ specific content of the thrown exception, for instance to chain active exceptions.

Otherwise, a catch block catching a foreign exception is allowed:

- to resume normal execution, thereby stopping propagation of the foreign exception and deleting it, or
- to re-throw the foreign exception. In that case, the original exception object must be unaltered by the C++ runtime.

A catch-all block may be executed during forced unwinding. For instance, a `longjmp` may execute code in a `catch(...)` during stack unwinding. However, if this happens, unwinding will proceed at the end of the catch-all block, whether or not there is an explicit re-throw.

Setting the low 4 bytes of exception class to C++\0 is reserved for use by C++ run-times compatible with the common C++ ABI.

Chapter 5

Conventions

1

¹This chapter is used to document some features special to the Intel386 ABI. The different sections might be moved to another place or removed completely.

5.1 C++

For the C++ ABI we will use the IA-64 C++ ABI and instantiate it appropriately.
The current draft of that ABI is available at:

<http://mentorembdedded.github.io/cxx-abi/>

Chapter 6

Alternate Code Sequences For Security

6.1 Code Sequences without PLT

Procedure Linkage Table (PLT) is used to access external functions defined in shared object and support

Lazy symbol resolution The function address is resolved only when it is called the first time at run-time.

Canonical function address The PLT entry of the external function is used as its address, aka function pointer.

The first instruction in the PLT entry is an indirect branch via the Global Offset Table (GOT) entry of the external function, which is set up in such a way that it will be updated to the address of the function body the first time when the function is called. Since the GOT entry is writable, any address may be written to it at run-time, which is a potential security risk.

6.1.1 Indirect Call via the GOT Slot

Different code sequences are used to avoid PLT when position independent code (PIC) is enabled and disabled:

Figure 6.1: Function Call without PLT (PIC)

<pre>extern void func (void); func ();</pre>	<pre>.globl func Load GOT base into reg call *func@GOT(%reg)</pre>
--	--

Either caller-save or callee-save registers can be used for GOT base to call an external function with PIC.

Figure 6.2: Function Call without PLT (Non-PIC)

<pre>extern void func (void); func ();</pre>	<pre>.globl func call *func@GOT</pre>
--	---

In both PIC and non-PIC cases, the direct branch is replaced by an indirect branch via the GOT slot, which is similar to the first instruction in the PLT slot.

Figure 6.3: Function Address without PLT (PIC)

<pre>extern void func (void); void* ptr (void) { return func; }</pre>	<pre>.globl func func: Load GOT base into eax movl func@GOT(%eax), %eax ret</pre>
---	---

Figure 6.4: Function Address without PLT (Non-PIC)

<pre>extern void func (void); void* ptr (void) { return func; }</pre>	<pre>.globl func func: movl func@GOT, %eax ret</pre>
---	--

Instead using the PLT slot as function address, the function address is retrieved from the GOT slot.

If linker determines the function is defined locally, it converts indirect branch via the GOT slot to direct branch with a `nop` prefix and converts load via the GOT slot to load immediate or `lea`, see Section A.2 for details.

After dynamic linker resolved all symbols by updating GOT entries with symbol addresses, GOT can be made read-only and overwriting GOT becomes a hard error immediately. Since PLT is no longer used to call external function, lazy symbol resolution is disabled and a function can only be interposed during symbol resolution at startup. Tools and features which depend on lazy symbol resolution will not work properly. However, there are also a few side benefits:

No extra direct branch to PLT entry Since indirect branch is 6 byte long and direct branch is 5 byte long, when indirect branch via the GOT slot is used to call a local function, code size will be increased by one byte for each call. Since one PLT slot has 16 bytes, there will be code size increase when indirect branch via the GOT slot is used to call an external function more than 16 times.

Custom calling convention Since external function is called directly via the GOT slot, instead of invoking dynamic linker to lookup function symbol when called the first time, parameters can be passed differently from what is specified in this document.

6.1.2 Thread-Local Storage without PLT

TLS code sequences for general and local dynamic models can be updated to replace direct call to `__tls_get_addr` via the PLT entry, with indirect call to `__tls_get_addr` via the GOT slot, see Figure 6.5. Since direct `call`

instruction is 4-byte long and indirect `call` instruction is 5-byte long, the extra one byte must be handled properly.

Figure 6.5: `__tls_get_addr` Call

Direct via PLT		Indirect via GOT	
<code>call</code>	<code>__tls_get_addr@PLT</code>	<code>call</code>	<code>*__tls_get_addr@GOT(%reg)</code>

General Dynamic Model for Global Variable

For general dynamic model, encoding of `leal` instruction before `call` instruction is changed from 7 bytes to 6 bytes to make room for indirect call:

```
extern __thread int x;
```

the following alternate code sequence loads address of `x` into `%eax` without PLT:

Table 6.1: General Dynamic Model Code Sequence

With PLT				Without PLT			
0x00	<code>leal</code>	<code>x@tls_gd(,%ebx,1), %eax</code>		0x00	<code>leal</code>	<code>x@tls_gd(%reg), %eax</code>	
0x07	<code>call</code>	<code>__tls_get_addr@PLT</code>		0x06	<code>call</code>	<code>*__tls_get_addr@GOT(%reg)</code>	

Either caller-save or callee-save registers can be used as GOT base for `R_386_TLS_GD` relocation against `x` and calling `__tls_get_addr`.

Static Thread-Local Variable

For local dynamic model, indirect call is used, instead of direct call:

```
static __thread int x;
```

the following alternate code sequence loads the address of the TLS block of the module, which contains variable `x`, into `%eax` without PLT:

Table 6.2: Local Dynamic Model Code Sequence

With PLT		Without PLT			
0x00	leal	x@tlsldm(%ebx), %eax	0x00	leal	x@tlsldm(%reg), %eax
0x06	call	____tls_get_addr@PLT	0x06	call	*____tls_get_addr@GOT(%reg)

As with general dynamic model, either caller-save or callee-save registers can be used as GOT base for R_386_TLS_LDM relocation against x and calling ____tls_get_addr.

TLS Linker Optimization

Since the code sequence with indirect call for general dynamic model has the same length as the one with direct call, linker just needs to recognize new instruction pattern to convert general dynamic access to initial exec or local exec accesses.

General Dynamic to Initial Exec To load address of x into %eax:

Table 6.3: GD -> IE Code Transition

GD		IE			
0x00	leal	x@tlsgd(%reg), %eax	0x00	movl	%gs:0, %eax
0x06	call	*____tls_get_addr@GOT(%reg)	0x06	subl	\$x@gottpoff(%reg), %eax

General Dynamic to Local Exec To load address of x into %eax:

Table 6.4: GD -> LE Code Transition

GD		LE			
0x00	leal	x@tlsgd(%reg), %eax	0x00	movl	%gs:0, %eax
0x06	call	*____tls_get_addr@GOT(%reg)	0x06	subl	\$x@tpoff, %eax

Local Dynamic to Local Exec For local dynamic model to local exec model transition, linker generates a 6-byte nop instruction, instead of a 1-byte nop

instruction plus a 4-byte `nop` instruction, after `mov` instruction, to account for the extra byte with indirect branch. To load the address of the TLS block of the module, which contains variable `x`, into `%eax` without PLT:

Table 6.5: LD -> LE Code Transition

LD			LE		
0x00	<code>leal</code>	<code>x@tlsldm(%reg), %eax</code>	0x00	<code>movl</code>	<code>%gs:0, %eax</code>
0x06	<code>call</code>	<code>*__tls_get_addr@GOT(%reg)</code>	0x06	<code>leal</code>	<code>0(%esi), %esi</code>

Chapter 7

Intel MPX Extension

Intel MPX (Memory Protection Extensions) provides 4 64-bit wide bound registers (`%bnd0 - %bnd3`). For purpose of function return, the lower 32 bits of `%bnd0` specify lower bound of function return, and the upper 32 bits specify upper bound of function return. The upper bound is represented in one's complement form.

7.1 Parameter Passing and Returning of Values

7.1.1 Bounds Passing

Intel MPX provides ISA extensions that allow passing bounds for a pointer argument that specify memory area that may be legally accessed by dereferencing the pointer. This paragraph describes how the bounds are passed to the callee.

Figure 7.1: Bound Register Usage

Register	Usage	Preserved across function calls
<code>%bnd0</code>	used to return bounds of pointer return value	No
<code>%bnd1–%bnd3</code>	scratch registers	No

Several functions used in the description below are defined as follows:

BOUND_MAP_STORE(bnd, addr, ptr) This function executes Intel MPX `bndstx` instruction. `ptr` argument is used to initialize index field of the memory operand of the `bndstx` instruction, `addr` is encoded in base and/or displacement fields of the memory operand, `bnd` is encoded in the register operand.

BOUND_MAP_LOAD(addr, ptr) This function executes Intel MPX `bndl dx` instruction. `ptr` argument is used to initialize index field of the memory operand of the `bndl dx` instruction, `addr` is encoded in base and/or displacement fields of the memory operand.

The bounds associated with each pointer contained in the fourbyte are passed in a CPU defined manner by executing `BOUND_MAP_STORE(bnd, addr, ptr)` function, where `bnd` is the current bounds of the pointer argument, `addr` is the address of the pointer argument's stack location, `ptr` is the actual value of the pointer argument. If the fourbyte may contain parts of partially overlapping pointers, then bounds associated with the pointers are ignored and special bounds that allow accessing all memory are passed for such pointers. The callee fetches the passed bounds using `BOUND_MAP_LOAD(addr, ptr)`, where `addr` is the same address passed to the corresponding `BOUND_MAP_STORE` in the caller, and `ptr` is the actual value of the pointer parameter fetched by the callee from a stack location.

When passing arguments with bounds to functions, function prototypes must be provided. Otherwise, the run-time behavior is undefined.

7.1.2 Returning of Bounds

The returning of bounds is done according to the following algorithm:

1. When the value is returned in memory, on return `%bnd0` must contain bounds of the “hidden” first argument that has been passed in by the caller.
2. When a pointer value is returned, on return `%bnd0` must contain bounds of the pointer value.

Chapter 8

Intel CET Extension

Intel CET (Control-flow Enforcement Technology) Extension includes:

IBT (Indirect Branch Tracking) Branch protection to defend against Jump/Call Oriented Programming.

SHSTK (Shadow Stack) Return address protection to defend against Return Oriented Programming,

8.1 Program Loading

8.1.1 Process `GNU_PROPERTY_X86_FEATURE_1_IBT`

On an IBT capable processor, the following steps should be taken:

1. When loading an executable without interpreter, enable IBT if `GNU_PROPERTY_X86_FEATURE_1_IBT` is set on the executable.
2. When loading an executable with an interpreter, enable IBT if `GNU_PROPERTY_X86_FEATURE_1_IBT` is set on the interpreter. The interpreter should disable IBT if `GNU_PROPERTY_X86_FEATURE_1_IBT` isn't set on the executable.
3. After IBT is enabled, when loading a shared object without `GNU_PROPERTY_X86_FEATURE_1_IBT`:

- (a) If legacy interwork is allowed, then mark all pages in executable `PT_LOAD` segments in legacy code page bitmap. Failure of legacy code page bitmap allocation causes an error.
- (b) If legacy interwork isn't allowed, it causes an error.

8.1.2 Process `GNU_PROPERTY_X86_FEATURE_1_SHSTK`

On a SHSTK capable processor, the following steps should be taken:

1. When loading an executable without interpreter, enable SHSTK if `GNU_PROPERTY_X86_FEATURE_1_SHSTK` is set on the executable.
2. When loading an executable with an interpreter, enable SHSTK if `GNU_PROPERTY_X86_FEATURE_1_SHSTK` is set on the interpreter. The interpreter should disable SHSTK if `GNU_PROPERTY_X86_FEATURE_1_SHSTK` isn't set or any shared objects loaded via the `DT_NEEDED` tag.
3. After SHSTK is enabled, when loading a shared object without `GNU_PROPERTY_X86_FEATURE_1_SHSTK`:
 - (a) If legacy interwork is allowed, SHSTK should be disabled.
 - (b) If legacy interwork isn't allowed, it causes an error.

8.2 Dynamic Linking

To support IBT, linker should generate a IBT-enabled PLT (Procedure Linkage Table) (see figure 8.1) together with a second PLT (see figure 8.2).

Figure 8.1: The IBT-enabled PLT

```
.PLT0: pushq    4(%ebx)           # GOT[1]
        jmp     *8(%ebx)        # GOT[2]
        nopl   0x0(%eax)
.PLT1: endbr32                   # 16 bytes from .PLT0
        pushq  $index1
        jmp     PLT0
        xchg  %ax,%ax
.PLT2: endbr32                   # 16 bytes from .PLT1
        pushq  $index2
        jmp     .PLT0
        xchg  %ax,%ax
.PLT3: ...
```

Figure 8.2: The Second IBT-enabled PLT

```
.SPLT1: endbr32
        jmp  *name1@GOT(%ebx)
        nopl 0x0(%eax,%eax,1)
.SPLT2: endbr32
        jmp  *name2@GOT(%ebx)           # 16 bytes from .SPLT1
        nopl 0x0(%eax,%eax,1)
.SPLT3: ...                             # 16 bytes from .SPLT2
```

Appendix A

Linker Optimization

This chapter describes optimizations which may be performed by linker.

A.1 Combine GOTPLT and GOT Slots

In the small and medium models, when there are both PLT and GOT references to the same function symbol, normally linker creates a GOTPLT slot for PLT entry and a GOT slot for GOT reference. A run-time JUMP_SLOT relocation is created to update the GOTPLT slot and a run-time GLOB_DAT relocation is created to update the GOT slot. Both JUMP_SLOT and GLOB_DAT relocations apply the same symbol value to GOTPLT and GOT slots, respectively, at run-time.

As an optimization, linker may combine GOTPLT and GOT slots into a single GOT slot and remove the run-time JUMP_SLOT relocation. It replaces the regular PLT entry:

Figure A.1: Procedure Linkage Table Entry Via GOTPLT Slot

```
.PLT: jmp      [GOTPLT slot]
      pushl   relocation index
      jmp     .PLT0
```

with an GOT PLT entry with an indirect jump via the GOT slot:

Figure A.2: Procedure Linkage Table Entry Via GOT Slot

```
.PLT: jmp      [GOT slot]
      nop
```

and resolves the PLT reference to the GOT PLT entry. Indirect `jmp` is an 5-byte instruction. `nop` can be encoded as a 3-byte instruction or a 11-byte instruction for 8-byte or 16-byte PLT slot. A separate PLT with 8-byte slots may be used for this optimization.

This optimization isn't applicable to the `STT_GNU_IFUNC` symbols since their GOTPLT slots are resolved to the selected implementation and their GOT slots are resolved to their PLT entries.

This optimization must be avoided if pointer equality is needed since the symbol value won't be cleared in this case and the dynamic linker won't update the GOT slot. Otherwise, the resulting binary will get into an infinite loop at run-time.

A.2 Optimize R_386_GOT32X Relocation

The Intel386 instruction encoding supports converting certain instructions on memory operand with `R_386_GOT32X` relocation against symbol, `foo`, into a different form on immediate operand if `foo` is defined locally.

Convert call, jmp and mov Convert memory operand of `call`, `jmp` and `mov` into immediate operand.

Table A.1: Call, Jmp and Mov Conversion

Memory Operand	Immediate Operand
<code>call *foo@GOT(%reg)</code>	<code>nop call foo</code>
<code>call *foo@GOT(%reg)</code>	<code>call foo nop</code>
<code>jmp *foo@GOT(%reg)</code>	<code>jmp foo nop</code>
<code>mov foo@GOT(%reg1), %reg2</code>	<code>lea foo@GOTOFF(%reg1), %reg2</code>

Convert Test and Binop Convert memory operand of `call`, `jmp`, `mov`, `test` and `binop` into immediate operand, where `binop` is one of `adc`, `add`, `and`, `cmp`, `or`, `sbb`, `sub`, `xor` instructions, when position-independent code is disabled.

Table A.2: Test and Binop Conversion

Memory Operand	Immediate Operand
<code>call *foo@GOT</code>	<code>nop call foo</code>
<code>call *foo@GOT</code>	<code>call foo nop</code>
<code>jmp *foo@GOT</code>	<code>jmp foo nop</code>
<code>mov foo@GOT, %reg</code>	<code>mov \$foo, %reg</code>
<code>test %reg, foo@GOT</code>	<code>test \$foo, %reg</code>
<code>binop foo@GOT, %reg</code>	<code>binop \$foo, %reg</code>
<code>call *foo@GOT(%reg)</code>	<code>nop call foo</code>
<code>call *foo@GOT(%reg)</code>	<code>call foo nop</code>
<code>jmp *foo@GOT(%reg)</code>	<code>jmp foo nop</code>
<code>mov foo@GOT(%reg1), %reg2</code>	<code>mov \$foo, %reg2</code>
<code>test %reg1, name@GOT(%reg2)</code>	<code>test \$foo, %reg1</code>
<code>binop name@GOT(%reg1), %reg2</code>	<code>binop \$foo, %reg2</code>

Index

`_UA_CLEANUP_PHASE`, 45
`_UA_FORCE_UNWIND`, 44
`_UA_SEARCH_PHASE`, 45
`_Unwind_Context`, 44
`_Unwind_DeleteException`, 43
`_Unwind_Exception`, 44
`_Unwind_ForcedUnwind`, 43, 44
`_Unwind_GetCFA`, 43
`_Unwind_GetGR`, 43
`_Unwind_GetIP`, 43
`_Unwind_GetLanguageSpecificData`, 43
`_Unwind_GetRegionStart`, 43
`_Unwind_RaiseException`, 43, 44
`_Unwind_Resume`, 43
`_Unwind_SetGR`, 43
`_Unwind_SetIP`, 43

auxiliary vector, 19

boolean, 8

byte, 6

C++, 60

Call Frame Information tables, 43

Convert call, jmp and mov, 73

Convert Test and Binop, 74

double quadword, 6

doubleword, 6

DWARF Debugging Information Format, 43

eightbyte, 6
exec, 17

fourbyte, 6

halfword, 6

longjmp, 44

Procedure Linkage Table, 34

quadword, 6

sixteenbyte, 6
size_t, 8

terminate(), 45
Thread-Local Storage, 37
twobyte, 6

Unwind Library interface, 43

word, 6