



C++ Application Binary Interface Standard for the ARM 64-bit Architecture

Document number: ARM IHI 0059B, current through AArch64 ABI release 1.0
Date of Issue: 22nd May 2013

Abstract

This document describes the C++ Application Binary Interface for the ARM 64-bit architecture.

Keywords

C++, Application Binary Interface, ABI, AArch64, C++ ABI, generic C++ ABI

How to find the latest release of this specification or report a defect in it

Please check the *ARM Information Center* (<http://infocenter.arm.com/>) for a later release if your copy is more than 3 months old (navigate to the *Software Development Tools* section, *Application Binary Interface for the ARM Architecture* subsection). Please report defects in this specification to *arm dot eabi at arm dot com*.

Licence

THE TERMS OF YOUR ROYALTY FREE LIMITED LICENCE TO USE THIS ABI SPECIFICATION ARE GIVEN IN SECTION 1.4, **Your licence to use this specification** (ARM contract reference **LEC-ELA-00081 V2.0**). PLEASE READ THEM CAREFULLY.

BY DOWNLOADING OR OTHERWISE USING THIS SPECIFICATION, YOU AGREE TO BE BOUND BY ALL OF ITS TERMS. IF YOU DO NOT AGREE TO THIS, DO NOT DOWNLOAD OR USE THIS SPECIFICATION.

THIS ABI SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES (SEE SECTION 1.4 FOR DETAILS).

Proprietary notice

ARM, Thumb, RealView, ARM7TDMI and ARM9TDMI are registered trademarks of ARM Limited. The ARM logo is a trademark of ARM Limited. ARM9, ARM926EJ-S, ARM946E-S, ARM1136J-S ARM1156T2F-S ARM1176JZ-S Cortex, and Neon are trademarks of ARM Limited. All other products or services mentioned herein may be trademarks of their respective owners.

Contents

1	ABOUT THIS DOCUMENT	3
1.1	Change Control	3
1.1.1	Current Status and Anticipated Changes	3
1.1.2	Change History	3
1.2	References	3
1.3	Terms and Abbreviations	3
1.4	Your Licence to Use This Specification	4
2	OVERVIEW	6
2.1	The Generic C++ ABI	6
2.2	The Exception Handling ABI for the ARM architecture	6
3	THE C++ ABI SUPPLEMENT	7
3.1	Summary of differences from and additions to the generic C++ ABI	7
3.2	Differences in detail	8
3.2.1	Representation of pointer to member function	8
3.2.2	Guard variables	8
3.2.3	Controlling Object Construction Order	8
3.2.4	ELF binding of static data guard variable symbols	8
3.2.5	Unwind Table Location	9
4	EH ABI LEVEL III: IMPLEMENTATION ABI FOR GNU LINUX	10
4.1	Introduction	10
4.2	Data Structures	10
4.3	Standard Runtime Initialization	10
4.4	Throwing an Exception	10
4.5	Catching an Exception	10
4.5.1	Overview of Catch Processing	10
4.5.2	The Personality Routine	10
4.5.3	Exception Handlers	10

1 ABOUT THIS DOCUMENT

1.1 Change Control

1.1.1 Current Status and Anticipated Changes

This document's status is released. Clarifications, extensions and minor changes should be expected.

1.1.2 Change History

Issue	Date	By	Change
00bet3	15 th December 2010	MGD	Beta release.
1.0	22 nd May 2013	RE	First public release.

1.2 References

This document refers to, or is referred to by, the following documents.

Ref	URL or other reference	Title
AAPCS64	IHI 0055	Procedure Call Standard for the ARM 64-bit Architecture
AAELF64	IHI 0056	ELF for the ARM 64-bit Architecture
CPPABI64	This document	C++ ABI for the ARM 64-bit Architecture
GC++ABI	http://mentorembdedded.github.io/cxx-abi/abi.html	Generic C++ ABI
Generic ELF	http://www.sco.com/developers/gabi/	Generic ELF, 17 th December 2003 Draft
ISO C++	ISO/IEC 14882:2003 (14882:1988 with <i>Technical Corrigendum</i>)	International Standard ISO/IEC 14882:2003 – Programming languages C++
LSB	http://refspecs.freestandards.org/LSB_4.0.0/LSB-Core-generic/LSB-Core-generic/book1.html	Linux Standards Base Core Specification 4.0
IA64EHABI	http://www.codesourcery.com/public/cxx-abi/abi-eh.html	Itanium C++ ABI: Exception Handling

1.3 Terms and Abbreviations

The *ABI for the ARM 64-bit Architecture* uses the following terms and abbreviations.

Term	Meaning
A32	The instruction set named <i>ARM</i> in the ARMv7 architecture; A32 uses 32-bit fixed-length instructions.
A64	The instruction set available when in AArch64 state.
AAPCS64	Procedure Call Standard for the ARM 64-bit Architecture (AArch64)
AArch32	The 32-bit general-purpose register width state of the ARMv8 architecture, broadly compatible with the ARMv7-A architecture.
AArch64	The 64-bit general-purpose register width state of the ARMv8 architecture.

Term	Meaning
ABI	Application Binary Interface: <ol style="list-style-type: none"> 1. The specifications to which an executable must conform in order to execute in a specific execution environment. For example, the <i>Linux ABI for the ARM Architecture</i>. 2. A particular aspect of the specifications to which independently produced relocatable files must conform in order to be statically linkable and executable. For example, the <i>C++ ABI for the ARM Architecture</i>, <i>ELF for the ARM Architecture</i>, ...
ARM-based	... based on the ARM architecture ...
Floating point	Depending on context floating point means or qualifies: (a) floating-point arithmetic conforming to IEEE 754 2008; (b) the ARMv8 floating point instruction set; (c) the register set shared by (b) and the ARMv8 SIMD instruction set.
Q-o-I	Quality of Implementation – a quality, behavior, functionality, or mechanism not required by this standard, but which might be provided by systems conforming to it. Q-o-I is often used to describe the tool-chain-specific means by which a standard requirement is met.
SIMD	Single Instruction Multiple Data – A term denoting or qualifying: (a) processing several data items in parallel under the control of one instruction; (b) the ARM v8 SIMD instruction set; (c) the register set shared by (b) and the ARMv8 floating point instruction set.
SIMD and floating point	The ARM architecture’s SIMD and Floating Point architecture comprising the floating point instruction set, the SIMD instruction set and the register set shared by them.
T32	The instruction set named Thumb in the ARMv7 architecture; T32 uses 16-bit and 32-bit instructions.

More specific terminology is defined when it is first used.

1.4 Your Licence to Use This Specification

IMPORTANT: THIS IS A LEGAL AGREEMENT (“LICENCE”) BETWEEN YOU (AN INDIVIDUAL OR SINGLE ENTITY WHO IS RECEIVING THIS DOCUMENT DIRECTLY FROM ARM LIMITED) (“LICENSEE”) AND ARM LIMITED (“ARM”) FOR THE SPECIFICATION DEFINED IMMEDIATELY BELOW. BY DOWNLOADING OR OTHERWISE USING IT, YOU AGREE TO BE BOUND BY ALL OF THE TERMS OF THIS LICENCE. IF YOU DO NOT AGREE TO THIS, DO NOT DOWNLOAD OR USE THIS SPECIFICATION.

“Specification” means, and is limited to, the version of the specification for the Applications Binary Interface for the ARM Architecture comprised in this document. Notwithstanding the foregoing, “Specification” shall not include (i) the implementation of other published specifications referenced in this Specification; (ii) any enabling technologies that may be necessary to make or use any product or portion thereof that complies with this Specification, but are not themselves expressly set forth in this Specification (e.g. compiler front ends, code generators, back ends, libraries or other compiler, assembler or linker technologies; validation or debug software or hardware; applications, operating system or driver software; RISC architecture; processor microarchitecture); (iii) maskworks and physical layouts of integrated circuit designs; or (iv) RTL or other high level representations of integrated circuit designs.

Use, copying or disclosure by the US Government is subject to the restrictions set out in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs (c)(1) and (2) of the Commercial Computer Software – Restricted Rights at 48 C.F.R. 52.227-19, as applicable.

This Specification is owned by ARM or its licensors and is protected by copyright laws and international copyright treaties as well as other intellectual property laws and treaties. The Specification is licensed not sold.

1. Subject to the provisions of Clauses 2 and 3, ARM hereby grants to LICENSEE, under any intellectual property that is (i) owned or freely licensable by ARM without payment to unaffiliated third parties and (ii) either embodied in the Specification or Necessary to copy or implement an applications binary interface compliant with this Specification, a perpetual, non-exclusive, non-transferable, fully paid, worldwide limited licence (without the right to sublicense) to use and copy this Specification solely for the purpose of developing, having developed, manufacturing, having manufactured, offering to sell, selling, supplying or otherwise distributing products which comply with the Specification.
2. THIS SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF SATISFACTORY QUALITY, MERCHANTABILITY, NONINFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE. THE SPECIFICATION MAY INCLUDE ERRORS. ARM RESERVES THE RIGHT TO INCORPORATE MODIFICATIONS TO THE SPECIFICATION IN LATER REVISIONS OF IT, AND TO MAKE IMPROVEMENTS OR CHANGES IN THE SPECIFICATION OR THE PRODUCTS OR TECHNOLOGIES DESCRIBED THEREIN AT ANY TIME.
3. This Licence shall immediately terminate and shall be unavailable to LICENSEE if LICENSEE or any party affiliated to LICENSEE asserts any patents against ARM, ARM affiliates, third parties who have a valid licence from ARM for the Specification, or any customers or distributors of any of them based upon a claim that a LICENSEE (or LICENSEE affiliate) patent is Necessary to implement the Specification. In this Licence; (i) "affiliate" means any entity controlling, controlled by or under common control with a party (in fact or in law, via voting securities, management control or otherwise) and "affiliated" shall be construed accordingly; (ii) "assert" means to allege infringement in legal or administrative proceedings, or proceedings before any other competent trade, arbitral or international authority; (iii) "Necessary" means with respect to any claims of any patent, those claims which, without the appropriate permission of the patent owner, will be infringed when implementing the Specification because no alternative, commercially reasonable, non-infringing way of implementing the Specification is known; and (iv) English law and the jurisdiction of the English courts shall apply to all aspects of this Licence, its interpretation and enforcement. The total liability of ARM and any of its suppliers and licensors under or in relation to this Licence shall be limited to the greater of the amount actually paid by LICENSEE for the Specification or US\$10.00. The limitations, exclusions and disclaimers in this Licence shall apply to the maximum extent allowed by applicable law.

ARM Contract reference LEC-ELA-00081 V2.0 AB/LS (9 March 2005)

2 OVERVIEW

The C++ ABI for the ARM 64-bit architecture (CPPABI64) comprises the following sub-components.

- The generic C++ ABI, summarized in §2.1, is the referenced base standard for this component.
- The *C++ ABI supplement* in §3 details ARM-specific additions to and deviations from the generic standard.
- The generic C++ exception handling ABI summarized in §2.2, describes the language-independent and C++-specific aspects of exception handling.
- The C++ exception handling ABI supplement for GNU/Linux in §4 details ARM-specific additions to and deviations from the generic standard for GNU/Linux systems.

The generic C++ ABI is implicitly an SVr4-based standard, and takes an SVr4 position on symbol visibility and vague linkage. This document does not cover extensions for DLL-based environment.

2.1 The Generic C++ ABI

The generic C++ ABI [GC++ABI] (originally developed for SVr4 on Itanium) specifies:

- The layout of C++ non-POD class types in terms of the layout of POD types (specified for *this* ABI by the *Procedure Call Standard for the ARM Architecture* [AAPCS64]).
- How class types requiring copy construction are passed as parameters and results.
- The content of run-time type information (RTTI).
- Necessary APIs for object construction and destruction.
- How names with linkage are mangled (name mangling).

The generic C++ ABI refers to a separate Itanium-specific specification of exception handling. When the generic C++ ABI is used as a component of *this* ABI, corresponding reference must be made to the generic C++ exception handling ABI as summarised in §2.2, and the C++ exception handling ABI supplement in §4.

2.2 The Exception Handling ABI for the ARM architecture

In common with [IA64EHABI], this ABI specifies table-based unwinding that separates language-independent unwinding from language specific aspects. The [IA64EHABI] specification describes three levels of ABI:

- Level I. A Base API, interfaces common to all languages and implementations.
- Level II. The C++ ABI, interfaces for interoperability of C++ implementations.
- Level III. The Implementation ABI specific to a particular runtime implementation.

The AArch64 C++ EH ABI uses Level I and Level II of the Itanium exception handling ABI as specified.

Section 4 describes the EH Level III ABI used on GNU/Linux systems.

3 THE C++ ABI SUPPLEMENT

3.1 Summary of differences from and additions to the generic C++ ABI

This section summarizes the differences between the *C++ ABI for the ARM 64-bit architecture* and the generic C++ ABI. Section numbers in captions refer to the generic C++ ABI specification. Larger differences are detailed in subsections of §3.2.

GC++ABI §2.2 POD Data Types

The GC++ABI defines the way in which empty class types are laid out. For the purposes of parameter passing in [AAPCS64], a parameter whose type is an empty class shall be treated as if its type were an aggregate with a single member of type unsigned byte.

Note Of course, the single member has undefined content.

GC++ABI §2.3 Member Pointers

The pointer to member function representation differs from that used by Itanium. See §3.2.1.

GC++ABI §2.8 Initialization guard variables

This ABI only specifies the bottom bit of the guard variable. See §3.2.2.

GC++ABI §3.1.4 Return Values

When a return value has a non-trivial copy constructor or destructor, the address of the caller-allocated temporary is passed in the Indirect Result Location Register (r8) rather than as an implicit first parameter before the `this` parameter and user parameters.

GC++ABI §3.3.4 Controlling Object Construction Order

Global object construction is managed in a simplified way under this ABI. See §3.2.3.

GC++ABI §3.4 Demangler

This ABI provides the demangler interface as specified in the generic C++ ABI. The ABI does not specify the format of the demangled string.

GC++ABI §5.2.2 Static Data

If a static datum and its guard variable are emitted in the same COMDAT group, the ELF binding [Generic ELF] for both symbols must be `STB_GLOBAL`, not `STB_WEAK` as specified in [GC++ABI64]. §3.2.4 justifies this requirement.

GC++ABI §5.3 Unwind Table Location

The exception unwind table shall be located by use of program header entries of type `PT_AARCH64_UNWIND`. See §3.2.5.

(No section in the generic C++ ABI) A library nothrow new function must not examine its 2nd argument

Library versions of the following functions must not examine their second argument.

```
::operator new(std::size_t, const std::nothrow_t&)
::operator new[](std::size_t, const std::nothrow_t&)
```

(The second argument conveys no useful information other than through its presence or absence, which is manifest in the mangling of the name of the function. This ABI therefore allows code generators to use a potentially invalid second argument – for example, whatever value happens to be in R1 – at a point of call).

(No section in the generic C++ ABI, but would be §2.2) POD data types

Pointers to extern "C++" functions and pointers to extern "C" functions are interchangeable if the function types are otherwise identical.

In order to be used by the library helper functions described below, implementations of constructor and destructor functions (complete, sub-object, deleting, and allocating) must have a type compatible with:

```
extern "C" void (*)(void* /* , other argument types if any */);
```

(No section in the generic C++ ABI) Namespace and mangling for the `va_list` type

The type `__va_list` is in namespace `std`. The type name of `va_list` therefore mangles to `St9__va_list`.

3.2 Differences in detail

3.2.1 Representation of pointer to member function

The generic C++ ABI [GC++ABI] specifies that a pointer to member function is a pair of words $\langle ptr, adj \rangle$. The least significant bit of ptr discriminates between (0) the address of a non-virtual member function and (1) the offset in the class's virtual table of the address of a virtual function.

This encoding cannot work for the AArch64 instruction set where the architecture reserves all bits of code addresses.

This ABI specifies that adj contains twice the $this$ adjustment, plus 1 if the member function is virtual. The least significant bit of adj then makes exactly the same discrimination as the least significant bit of ptr does for Itanium.

A pointer to member function is NULL when $ptr = 0$ and the least significant bit of adj is zero.

3.2.2 Guard variables

The generic C++ ABI [GC++ABI] specifies the bottom byte of a static variable guard variable shall be 0 when the variable is not initialized, and 1 when it is. All other bytes are platform defined.

This ABI instead only specifies the value bit 0 of the static guard variable; all other bits are platform defined. Bit 0 shall be 0 when the variable is not initialized and 1 when it is.

3.2.3 Controlling Object Construction Order

The generic ABI specifies a `#pragma` and `.priority_init` section type to allow the user to specify object construction order. This scheme is not in wide use and so this ABI uses a different scheme which has several pre-existing implementations.

The compiler is responsible for sequencing the construction of top-level static objects defined in a translation unit in accordance with the requirements of the C++ standard. The run-time environment (helper-function library) sequences the initialization of one translation unit after another. The global constructor vector provides the interface between these agents as follows:

- Each translation unit provides a fragment of the constructor vector in an ELF section called `.init_array` of type `SHT_INIT_ARRAY` (=0xE) and section flags `SHF_ALLOC + SHF_WRITE`.
- Each element of the vector contains the address of a function of type `extern "C" void (* const)(void)` that, when called, performs part or all of the global object construction for the translation unit. Producers must treat `.init_array` sections as if they were read-only.
- The appropriate entry for an element referring to, say, `__sti_file` that constructs the global static objects in `filecpp`, is 0 relocated by `R_AARCH64_ABS64` (`__sti_file`).
- Object construction order may be controlled by appending an unsigned integer in the range 0-65535 (formatted as if by `printf("%05d", priority)`) to the name of the section. The linker must lay these sections out in ascending lexicographical order.
- Sections without a priority number appended are assumed to have a lower priority than those sections with a priority number. The linker should lay out sections without a priority number after those sections with.
- The priority values 0 to 100 inclusive are reserved to the implementation.
- Run-time support code iterates through the global constructor vector in increasing address order calling each identified initialization function in order.

3.2.4 ELF binding of static data guard variable symbols

The generic C++ standard [GC++ABI] states at the end of §5.2.2:

Local static data objects generally have associated guard variables used to ensure that they are initialized only once (see 3.3.2). If the object is emitted using a COMDAT group, the guard variable must be too. It is

suggested that it be emitted in the same COMDAT group as the associated data object, but it may be emitted in its own COMDAT group, identified by its name. In either case, it must be weak.

In effect the generic standard permits a producer to generate one of two alternative structures. Either:

```
COMDAT Group (Variable Name) {
    Defines Variable Name      // ELF binding STB_GLOBAL, mangled name
    Defines Guard Variable Name // ELF binding STB_WEAK, mangled name ...
}
// (... this ABI requires STB_GLOBAL binding)
```

Or:

```
COMDAT Group (Variable Name) {
    Defines Variable Name      // ELF binding STB_GLOBAL, mangled name
}
+
COMDAT Group (Guard Variable Name) {
    Defines Guard Variable Name // ELF binding STB_WEAK, mangled name
}
```

A link step involving multiple groups of the first kind causes no difficulties. A linker must retain only one copy of the group and there will be one definition of Variable Name and one weak definition of Guard Variable Name.

A link step involving pairs of groups of the second kind also causes no difficulties. A linker must retain one copy of each group so there will be one definition of Variable Name and one weak definition of Guard Variable Name.

A link step involving a group of the first kind and a pair of groups of the second kind generates two sub-cases.

- If the linker discards the group that defines two symbols there is no problem.
- If the linker retains the group that defines both Variable Name and Guard Variable Name it must nonetheless retain the group called Guard Variable Name. There are now two definitions of Guard Variable Name with ELF binding STB_WEAK.

In this second case there is no problem provided the linker picks one of the definitions.

Unfortunately, [Generic ELF] does not specify how linkers must process multiple weak definitions when there is no non-weak definition to override them. If a linker faults duplicate weak definitions there will be a functional failure.

This ABI requires the ELF binding of Guard Variable Name in the first structure to be STB_GLOBAL.

The rules codified in [Generic ELF] then make all three linking scenarios well defined and it becomes possible to link the output of compilers such as armcc that choose the first structure with the output of those such as gcc that choose the second without relying on linker behavior that the generic ELF standard leaves unspecified.

3.2.5 Unwind Table Location

Exception tables are located in sections with the name `.eh_frame` and `.eh_frame_hdr`. Linkers shall put the `.eh_frame_hdr` section in a single text segment, with a `PT_AARCH64_UNWIND` program table entry identifying the unwind table header location.

4 EH ABI LEVEL III: IMPLEMENTATION ABI FOR GNU LINUX

4.1 Introduction

This section describes the Exception Handling Implementation ABI for GNU Linux systems.

It specifies:

- The format of the unwind tables
- Standard Runtime Initialization features
- Throwing an Exception
- Catching an Exception

This section follows the layout of [IA64EHABI] §3.

4.2 Data Structures

The format of the exception tables is as specified in [LSB] §II.11.6 (Exception Frames).

The codes used to describe the encoding of pointers used in the exception frame tables, are the values described in [LSB] §II.11.5.1 (DWARF Exception Header Encoding).

Note that in particular that the layout of the Language Specific Data Area (LSDA) is not specified by this ABI. The structure and layout of a LSDA is specific to a particular implementation of a personality routine.

4.3 Standard Runtime Initialization

See [IA64EHABI] §3.3.

4.4 Throwing an Exception

See [IA64EHABI] §3.4.

4.5 Catching an Exception

4.5.1 Overview of Catch Processing

Stack unwinding itself is begun by calling `__Unwind_RaiseException()`, and performed by the unwind library. See [IA64EHABI] §3.5 for a summary.

4.5.2 The Personality Routine

The personality routine is specified in [IA64EHABI] §2.5.2.

4.5.3 Exception Handlers

The behavior of exception handlers is described in [IA64EHABI] §2.5.3.