# LoongArch Toolchain Conventions

Loongson Technology Corporation Limited

Version 1.00

## **Table of Contents**

Compiler Options	2
Rationale	2
Configuring Target ISA	3
Configuring Target ABI	4
GNU Target Triplets and Multiarch Specifiers	6
C/C++ Preprocessor Built-in Macro Definitions	8

Note: In this document, the terms "architecture", "instruction set architecture" and "ISA" are used synonymously and refer to a certain set of instructions and the set of registers they can operate upon.

## **Compiler Options**

#### Rationale

Compiler options that are specific to LoongArch should denote a change in the following compiler settings:

- Target architecture: the allowed set of instructions and registers to be used by the compiler.
- Target ABI type: the data model and calling conventions.
- Target microarchitecture: microarchitectural features that guides compiler optimizations.

For this model, two categories of LoongArch-specific compiler options should be implemented :

- **Basic options**: select the base configuration of the compilation target. (include only -march -mabi and -mtune)
- Extended options: make incremental changes to the target configuration.

Table 1. Basic Options

Option	Possible values	Description
-march=	native loongarch64 la464	Selects the target architecture, i.e. the basic set of ISA modules to be enabled.
-mabi=	lp64dlp64flp64silp32d ilp32filp32s	Selects the base ABI type.
-mtune=	native loongarch64 la464	Selects the type of target microarchitecture, defaults to the value of -march. The set of possible values should be a superset of -march values.

• Valid parameter values of -march and -mtune options should correspond to actual LoongArch processor implementations / families.

• In principle, different -march values should not imply the same set of ISA modules.

Table	2.	Extended	Options
rubic	∠.	LATCHUCU	options

Option	Possible values	Description
-msoft -float		Prevent the compiler from generating hardware floating-point instructions, and adjust the selected base ABI type to use soft- float calling convention. (The adjusted base ABI identifier should have suffix s.)
-msingl e-float		Allow generating 32-bit floating-point instructions, and adjust the selected base ABI type to use 32-bit FP calling convention. (The adjusted base ABI identifier should have suffix f.)
-mdoubl e-float		Allow generating 32- and 64-bit floating-point instructions. and adjust the selected base ABI type to use 64-bit FP calling convention. (The adjusted base ABI identifier should have suffix d.)
-mfpu=	64 32 θ none (equivalent to θ)	Selects the allowed set of basic floating-point instructions and registers. This option should not change the FP calling convention unless it's necessary. (The implementation of this option is not mandatory. It is recommended to use -m*-float options in software projects and scripts.)

As a general rule, the effect of all LoongArch-specific compiler options that are given for one compiler invocation should be as if they are processed in the order they appear on the command line. The only exception to this rule is  $-m^*-float$ : their configuration of floating-point instruction set and calling convention will not be changed by subsequent options other than  $-m^*-float$ .

### **Configuring Target ISA**

The LoongArch ISA is organized in a "base-extension" manner. For future updates, each component in the *base* or *extened* part of the ISA may evolve independently while keeping compatibility with previous versions of itself.

For this purpose, the compiler should make a modular abstraction about the target ISA. The ISA modules are divided into two categories : **base architectures** and **ISA extensions**. A **base architecture** is the core component of the target ISA, which defines the *base* set of functionalities like integer and floating-point operations, and is decided by the value of -march. An **ISA extension** may represent either the base of a certain *extended* ISA component or an incremental update to it, and is enabled / disabled by extended options.

For a complier, the final ISA configuration should be derived by applying various ISA extension configurations from extended options to the base ISA, which is selected via the **-march** option and consists of the base architecture and the set of default ISA extensions.



Among all ISA modules listed below, the compiler should at least implement one base architecture.

Table 3. Base Architecture

Name	-march values that selects this module	Description
LA64 basic architecture v1.00 (1a64v100)	loongarch64 la464	ISA defined in LoongArch Reference Manual - Volume 1: Basic Architecture v1.00.

The following table lists all ISA extensions that should be abstracted by the compiler and the compiler

Table 4. ISA extensions

Name	Related compiler options	Description
Basic Floating- Point Processing Unit (fpu*)	-mfpu=* (Possible values of *: none 32 64)	Selects the allowed set of basic floating-point instructions and floating-point registers. This is a constituent part of the base architecture, where it gets its default value.

The following table lists the properties of all target CPU models that can serve as arguments to **-march** and **-mtune** options at the same time.

Table 5. Target CPU Models (-march=<model> and -mtune=<model>)

Name / Value	ISA modules that are enabled by default	Target of optimization
native	auto-detected with <a href="cpucfg">cpucfg</a> (only on native and cross-native compilers)	auto-detected microarchitecture model / features
loongarch64	la64v100[fpu64]	generic LoongArch LA64 processors
la464	la64v100[fpu64]	LA464 processor core

#### **Configuring Target ABI**

Like configuring the target ISA, a complete ABI configuration of LoongArch consists of two parts, the **base ABI** and the **ABI extension**. The former describes the data model and calling convention in general, while the latter denotes an overall adjustment to the base ABI, which may require support from certain ISA extensions.

Please be noted that there is only ONE ABI extension slot in an ABI configuration. They do not combine with one another, and are, in principle, mutually incompatible.

A new ABI extension type will not be added to this document unless it implies certain significant performance / functional advantage that no compiler optimization techniques can provide without altering the ABI.

There are six base ABI types, whose standard names are the same as the -mabi values that select them. The compiler may choose to implement one or more of these base ABI types, possibly according to the range of implemented target ISA variants.

Standard name	Data model	Bit-width of argument / return value GPRs / FPRs
lp64d	LP64	64 / 64
lp64f	LP64	64 / 32
lp64s	LP64	64 / (none)
ilp32d	ILP32	32 / 64
ilp32f	ILP32	32 / 32
ilp32s	ILP32	32 / (none)

Table 6. Base ABI Types

The following table lists all ABI extension types and related compiler options. A compiler may choose to implement any subset of these extensions that contains base.

The default ABI extension type is base when referring to an ABI type with only the "base" component.

Table 7. ABI Extension Types

Name	Compiler options	Description
base	(none)	conforms to the LoongArch ELF psABI

The compiler should know the default ABI to use during its build time. If the ABI extension type is not explicitly configured, base should be used.

In principle, the target ISA configuration should not affect the decision of the target ABI. When certain ISA feature required by explicit (i.e. from the compiler's command-line arguments) ABI configuration cannot be met due constraints imposed by ISA options, the compiler should abort with an error message to complain about the conflict.

When the ABI is not fully constrained by the compiler options, the default configuration of either the base ABI or the ABI extension, whichever is missing from the command line, should be attempted. If this default ABI setting cannot be implemented by the explicitly configured target ISA, the expected behavior is **undefined** since the user is encouraged to specify which ABI to use when choosing a smaller instruction set than the default.

In this case, it is suggested that the compiler should abort with an error message, however, for userfriendliness, it may also choose to ignore the default base ABI or ABI extension and select a viable fallback ABI for the currently enabled ISA modules with caution. It is also recommended that the compiler should notify the user about the ABI change, optionally with a compiler warning. For example, passing -mfpu=none as the only command-line argument may cause a compiler configured with 1p64d / base default ABI to automatically select 1p64s / base instead.

When the target ISA configuration cannot be uniquely decided from the given compiler options, the buildtime default should be consulted first. If the default ISA setting is insufficient for implementing the ABI configuration, the compiler should try enabling the missing ISA modules according to the following table, as long as they are not explicitly disabled or excluded from usage.

Base ABI type	ABI extension type	Minimal required ISA modules
lp64d	base	la64v100[fpu64]
lp64f	base	la64v100 fpu32
lp64s	base	la64v100 fpunone

Table 8. Minimal architecture requirements for implementing each ABI type.

## **GNU Target Triplets and Multiarch Specifiers**

**Target triplet** is a core concept in the GNU build system. It describes a platform on which the code runs and mostly consists of three fields: the CPU family / model (machine), the vendor (vendor), and the operating system name (os).

**Multiarch architecture apecifiers** are essentially standard directory names where libraries are installed on a multiarch-flavored filesystem. These strings are normalized GNU target triplets. See debian documentation for details.

This document recognizes the following machine strings for the GNU triplets of LoongArch:

Table 9. LoongArch machine strings :

machine	Description
loongarch64	LA64 base architecture (implies <a>1p64*</a> ABI)
loongarch32	LA32 base architecture (implies <pre>ilp32* ABI)</pre>

As standard library directory names, the canonical multiarch architecture specifiers of LoongArch should contain information about the ABI type of the libraries that are meant to be released in the binary form and installed there.

While the integer base ABI is implied by the machine field, the floating-point base ABI and the ABI extension type are encoded with two string suffices (<fabi-suffix><abiext-suffix>) to the os field of the specifier, respectively.

#### Table 10. List of possible <fabi-suffix>

<fabi-suffix></fabi-suffix>	Description
(empty string)	The base ABI uses 64-bit FPRs for parameter passing. (1p64d)
f32	The base ABI uses 32-bit FPRs for parameter passing. (1p64f)
sf	The base ABI uses no FPR for parameter passing. (1p64s)

Table 11. List of possible <abiext-suffix>

<abiext-suffix></abiext-suffix>	ABI extension type
(empty string)	base

(Note: Since in principle, The default ISA configuration of the ABI should be used in this binary-release scenario, it is not necessary to reserve multiple multiarch specifiers for one OS / ABI combination.)

Table	12.	List	of l	LoonaArch	mulitarch	specifiers
rubic	12.	LIGU	011	200119/11/011	mantaron	specificio

ABI type (Base ABI / ABI extension)	C Library	Kernel	Multiarch specifier
lp64d/base	glibc	Linux	loongarch64-linux- gnu
lp64f/base	glibc	Linux	loongarch64-linux- gnuf32
lp64s/base	glibc	Linux	loongarch64-linux- gnusf

ABI type (Base ABI / ABI extension)	C Library	Kernel	Multiarch specifier
lp64d/base	musl libc	Linux	loongarch64-linux- musl
lp64f/base	musl libc	Linux	loongarch64-linux- muslf32
lp64s/base	musl libc	Linux	loongarch64-linux- muslsf

## **C/C++ Preprocessor Built-in Macro Definitions**

The definitions listed below is not specific to LoongArch. Amount of LoongArch-specific code can be minimized by utilizing them, while achieving expected portability in most of cases.

Name	Possible Values	Description
BYTE_ORD ER	(omitted)	Byte order
FLOAT_WO RD_ORDER	(omitted)	Byte order for floating-point data
LP64 _LP64	(omitted)	Whether the ABI passes arguments in 64-bit GPRs and uses the LP64 data model
SIZEOF_S HORT	(omitted)	Width of C/C++ <pre>short</pre> type, in bytes
SIZEOF_I NT	(omitted)	Width of C/C++ int type, in bytes
SIZEOF_L ONG	(omitted)	Width of C/C++ long type, in bytes
SIZEOF_L ONG_LONG	(omitted)	Width of C/C++ long long type, in bytes
SIZEOF_I NT128	(omitted)	Width of C/C++int128 type, in bytes
SIZEOF_P OINTER	(omitted)	Width of C/C++ pointer types, in bytes
SIZEOF_P TRDIFF_T	(omitted)	Width of C/C++ <pre>ptrdiff_t</pre> type, in bytes
SIZEOF_S IZE_T	(omitted)	Width of C/C++ <pre>size_t</pre> type, in bytes
SIZEOF_W INT_T	(omitted)	Width of C/C++ <pre>wint_t</pre> type, in bytes
SIZEOF_W CHAR_T	(omitted)	Width of C/C++ <pre>wchar_t</pre> type, in bytes
SIZEOF_F LOAT	(omitted)	Width of C/C++ float type, in bytes
SIZEOF_D OUBLE	(omitted)	Width of C/C++ double type, in bytes
SIZEOF_L ONG_DOUBLE	(omitted)	Width of C/C++ long double type, in bytes

Table 13. Non-LoongArch-specific C/C++ Built-in Macros :

Apart from the generic definitions described above, some architecture-specific macros are still needed to convey those information strongly tied to the architecture; these macros are listed below.

Table 14. LoongArch-specific C/C++ Built-in Macros :

Name	Possible Values	Description	
loongarc h	1	Defined if the target is LoongArch.	
loongarc h_grlen	64	Bit-width of general purpose registers.	
loongarc h_frlen	0 32 64	Bit-width of floating-point registers (0 if there is no FPU).	
loongarc h_arch	"loongarch64" "la464"	Processor model as specified by -march. If -march is not present, the build-time default should be used.	
loongarc h_tune	"loongarch64" "la464"	Processor model as specified by -mtune. If -mtune is not present, the build-time default should be used.	
loongarc h_lp64	undefined or 1	Defined if ABI uses the LP64 data model and 64-bit GPRs for parameter passing.	
loongarc h_hard_flo at	undefined or 1	Defined if floating-point/extended ABI type is single or double.	
loongarc h_soft_flo at	undefined or 1	Defined if floating-point/extended ABI type is <b>soft</b> .	
loongarc h_single_f loat	undefined or 1	Defined if floating-point/extended ABI type is single.	
loongarc h_double_f loat	undefined or 1	Defined if floating-point/extended ABI type is double.	

For historical reasons, the earliest LoongArch C/C++ compilers provided some MIPS-style built-in macros. Because legacy code dependent on those macros is possibly still in use, compilers conformant to this specification may provide the macros as listed below.

Because the naming style and usage of these macros are more-or-less inconsistent with the other macros described above, there is learning cost involved in using these macros. As they bring no advantage over the other macros, it is not recommended for newer compilers to implement them; portable code should not assume existence of these macros, nor use them.

Table 15. C/C++	Built-in Macros	Provided fo	r Compatibility	with Historical Code
-----------------	-----------------	-------------	-----------------	----------------------

Name	Equivalent to	Description
loongarc h64	loongarch_grlen == 64	Similar to mips64; defined iff loongarch_grlen == 64.
_LOONGARCH _ARCH	loongarch_arch	n/a

Name	Equivalent to	Description
_LOONGARCH _TUNE	loongarch_tune	n/a
_LOONGARCH _SIM	n/a	Similar to _MIPS_SIM on MIPS; possible values are _ABILP64 (in case data model is LP64) and _ABILP32 (in case data model is ILP32; notice the omission of letter I).
_LOONGARCH _SZINT	SIZEOF_INT multiplied by 8	n/a
_LOONGARCH _SZLONG	SIZEOF_LONG multiplied by 8	n/a
_LOONGARCH _SZPTR	SIZEOF_POINTER multiplied by 8	n/a