## Introduction

This book is the primary reference for the Rust programming language.

[!NOTE] For known bugs and omissions in this book, see our <u>GitHub issues</u>. If you see a case where the compiler behavior and the text here do not agree, file an issue so we can think about which is correct.

#### **Rust releases**

Rust has a new language release every six weeks. The first stable release of the language was Rust 1.0.0, followed by Rust 1.1.0 and so on. Tools (rustc, cargo, etc.) and documentation (<u>Standard library</u>, this book, etc.) are released with the language release.

The latest release of this book, matching the latest Rust version, can always be found at <u>https://doc.rust-lang.org/reference/</u>. Prior versions can be found by adding the Rust version before the "reference" directory. For example, the Reference for Rust 1.49.0 is located at <u>https://doc.rust-lang.org/1.49.0/reference/</u>.

#### What The Reference is not

This book does not serve as an introduction to the language. Background familiarity with the language is assumed. A separate <u>book</u> is available to help acquire such background familiarity.

This book also does not serve as a reference to the <u>standard library</u> included in the language distribution. Those libraries are documented separately by extracting documentation attributes from their source code. Many of the features that one might expect to be language features are library features in Rust, so what you're looking for may be there, not here.

Similarly, this book does not usually document the specifics of rustc as a tool or of Cargo. rustc has its own <u>book</u>. Cargo has a <u>book</u> that contains a <u>reference</u>. There are a few pages such as <u>linkage</u> that still describe how rustc works.

This book also only serves as a reference to what is available in stable Rust. For unstable features being worked on, see the <u>Unstable Book</u>.

Rust compilers, including rustc, will perform optimizations. The reference does not specify what optimizations are allowed or disallowed. Instead, think of the compiled program as a black box. You can only probe by running it, feeding it input and observing its output. Everything that happens that way must conform to what the reference says.

#### How to use this book

This book does not assume you are reading this book sequentially. Each chapter generally can be read standalone, but will cross-link to other chapters for facets of the language they refer to, but do not discuss.

There are two main ways to read this document.

The first is to answer a specific question. If you know which chapter answers that question, you can jump to that chapter in the table of contents. Otherwise, you can press **s** or click the magnifying glass on the top bar to search for keywords related to your question. For example, say you wanted to know when a temporary value created in a let statement is dropped. If you didn't already know that the <u>lifetime of temporaries</u> is defined in the <u>expressions chapter</u>, you could search "temporary let" and the first search result will take you to that section.

The second is to generally improve your knowledge of a facet of the language. In that case, just browse the table of contents until you see something you want to know more about, and just start reading. If a link looks interesting, click it, and read about that section.

That said, there is no wrong way to read this book. Read it however you feel helps you best.

#### Conventions

Like all technical books, this book has certain conventions in how it displays information. These conventions are documented here.

• Statements that define a term contain that term in *italics*. Whenever that term is used outside of that chapter, it is usually a link to the section that has this definition.

An *example term* is an example of a term being defined.

• The main text describes the latest stable edition. Differences to previous editions are separated in edition blocks:

[!EDITION-2018] Before the 2018 edition, the behavior was this. As of the 2018 edition, the behavior is that.

• Notes that contain useful information about the state of the book or point out useful, but mostly out of scope, information are in note blocks.

[!NOTE] This is an example note.

• Example blocks show an example that demonstrates some rule or points out some interesting aspect. Some examples may have hidden lines which can be viewed by clicking the eye icon that appears when hovering or tapping the example.

```
[!EXAMPLE] This is a code example.
println!("hello world");
```

• Warnings that show unsound behavior in the language or possibly confusing interactions of language features are in a special warning box.

```
[!WARNING] This is an example warning.
```

• Code snippets inline in the text are inside <code> tags.

Longer code examples are in a syntax highlighted box that has controls for copying, executing, and showing hidden lines in the top right corner.

```
# // This is a hidden line.
fn main() {
    println!("This is a code example");
}
```

All examples are written for the latest edition unless otherwise stated.

• The grammar and lexical productions are described in the <u>Notation</u> chapter.

r[example.rule.label]

• Rule identifiers appear before each language rule enclosed in square brackets. These identifiers provide a way to refer to and link to a specific rule in the language (<u>e.g.</u>). The rule identifier uses periods to separate sections from most general to most specific

([destructors.scope.nesting.function-body] for example). On narrow screens, the rule name will collapse to display [\*].

The rule name can be clicked to link to that rule.

[!WARNING] The organization of the rules is currently in flux. For the time being, these identifier names are not stable between releases, and links to these rules may fail if they are changed. We intend to stabilize these once the organization has settled so that links to the rule names will not break between releases.

• Rules that have associated tests will include a Tests link below them (on narrow screens, the link is [T]). Clicking the link will pop up a list of tests, which can be clicked to view the test. For example, see [input.encoding.utf8].

Linking rules to tests is an ongoing effort. See the <u>Test summary</u> chapter for an overview.

#### Contributing

We welcome contributions of all kinds.

You can contribute to this book by opening an issue or sending a pull request to the Rust Reference repository. If this book does not answer your question, and you think its answer is in scope of it, please do not hesitate to file an issue or ask about it in the t-lang/doc stream on Zulip. Knowing what people use this book for the most helps direct our attention to making those sections the best that they can be. And of course, if you see anything that is wrong or is non-normative but not specifically called out as such, please also file an issue.

## Notation

#### Grammar

The following notations are used by the *Lexer* and *Syntax* grammar snippets:

Notation	Examples	Meaning
CAPITAL	KW_IF, INTEGER_LITERAL	A token produced by the lexer
ItalicCamelCase	LetStatement, Item	A syntactical production
string	x, while, *	The exact character(s)
x?	pub <sup>?</sup>	An optional item
x*	OuterAttribute <sup>*</sup>	0 or more of x
x <sup>+</sup>	MacroMatch <sup>+</sup>	1 or more of x
x <sup>ab</sup>	HEX_DIGIT <sup>16</sup>	a to b repetitions of x
Rule1 Rule2	fn Name Parameters	Sequence of rules in order
	u8   u16, Block   Item	Either one or another
[]	[b B]	Any of the characters listed

Notation	Examples	Meaning
[-]	[ a - z ]	Any of the characters in the range
~[]	~[b B]	Any characters, except those listed
~string	~ \n, ~ */	Any characters, except this sequence
()	(, Parameter) <sup>?</sup>	Groups items
U+xxxx	U+0060	A single unicode character
<text></text>	<any ascii="" char="" cr="" except=""></any>	An English description of what should be matched
Rule <sub>suffix</sub>	IDENTIFIER_OR_KEYWORD except crate	A modification to the previous rule

Sequences have a higher precedence than | alternation.

### String table productions

Some rules in the grammar — notably <u>unary operators</u>, <u>binary operators</u>, and <u>keywords</u> — are given in a simplified form: as a listing of printable strings. These cases form a subset of the rules regarding the <u>token</u> rule, and are assumed to be the result of a lexical-analysis phase feeding the parser, driven by a <u>DFA</u>, operating over the disjunction of all such string table entries.

When such a string in **monospace** font occurs inside the grammar, it is an implicit reference to a single member of such a string table production. See <u>tokens</u> for more information.

### **Grammar visualizations**

Below each grammar block is a button to toggle the display of a <u>syntax</u> <u>diagram</u>. A square element is a non-terminal rule, and a rounded rectangle is a terminal.

## **Common productions**

The following are common definitions used in the grammar. r[input.syntax] @root CHAR -> <a Unicode scalar value>

NUL -> U+0000

TAB -> U+0009

LF -> U+000A

CR -> U+000D

# Lexical structure

r[input]

## **Input format**

r[input.intro] This chapter describes how a source file is interpreted as a sequence of tokens.

See <u>Crates and source files</u> for a description of how programs are organised into files.

r[input.encoding]

### **Source encoding**

r[input.encoding.utf8] Each source file is interpreted as a sequence of Unicode characters encoded in UTF-8.

r[input.encoding.invalid] It is an error if the file is not valid UTF-8.

r[input.byte-order-mark]

## Byte order mark removal

If the first character in the sequence is U+FEFF (<u>BYTE ORDER</u> <u>MARK</u>), it is removed.

r[input.crlf]

### **CRLF** normalization

Each pair of characters U+000D (CR) immediately followed by U+000A (LF) is replaced by a single U+000A (LF).

Other occurrences of the character U+000D (CR) are left in place (they are treated as <u>whitespace</u>).

r[input.shebang]

### **Shebang removal**

r[input.shebang.intro] If the remaining sequence begins with the characters #!, the characters up to and including the first U+000A (LF) are removed from the sequence.

For example, the first line of the following file would be ignored:
#!/usr/bin/env rustx

```
fn main() {
    println!("Hello!");
}
```

}

r[input.shebang.inner-attribute] As an exception, if the #! characters are followed (ignoring intervening <u>comments</u> or <u>whitespace</u>) by a [ token, nothing is removed. This prevents an <u>inner attribute</u> at the start of a source file being removed.

[!NOTE] The standard library [include!] macro applies byte order mark removal, CRLF normalization, and shebang removal to the file it reads. The [include\_str!] and [include\_bytes!] macros do not.

r[input.tokenization]

### Tokenization

The resulting sequence of characters is then converted into tokens as described in the remainder of this chapter.

r[lex.keywords]

# Keywords

Rust divides keywords into three categories:

- <u>strict</u>
- <u>reserved</u>
- <u>weak</u>

r[lex.keywords.strict]

## Strict keywords

r[lex.keywords.strict.intro] These keywords can only be used in their correct contexts. They cannot be used as the names of:

- <u>Items</u>
- <u>Variables</u> and function parameters
- Fields and <u>variants</u>
- <u>Type parameters</u>
- Lifetime parameters or <u>loop labels</u>
- <u>Macros</u> or <u>attributes</u>
- <u>Macro placeholders</u>
- <u>Crates</u>

r[lex.keywords.strict.list] The following keywords are in all editions:

- as
- break
- const
- continue
- crate
- else
- enum
- extern
- false
- fn
- for
- if
- impl
- in
- let
- loop
- match
- mod

- move
- mut
- pub
- ref
- return
- self
- Self
- static
- struct
- super
- trait
- true
- type
- unsafe
- use
- where
- while

r[lex.keywords.strict.edition2018] The following keywords were added beginning in the 2018 edition.

- async
- await
- dyn

r[lex.keywords.reserved]

### **Reserved keywords**

r[lex.keywords.reserved.intro] These keywords aren't used yet, but they are reserved for future use. They have the same restrictions as strict keywords. The reasoning behind this is to make current programs forward compatible with future versions of Rust by forbidding them to use these keywords.

r[lex.keywords.reserved.list]

- abstract
- become
- box
- do
- final
- macro
- override
- priv
- typeof
- unsized
- virtual
- yield

r[lex.keywords.reserved.edition2018] The following keywords are reserved beginning in the 2018 edition.

• try

r[lex.keywords.reserved.edition2024] The following keywords are reserved beginning in the 2024 edition.

• gen

r[lex.keywords.weak]

## Weak keywords

r[lex.keywords.weak.intro] These keywords have special meaning only in certain contexts. For example, it is possible to declare a variable or method with the name union.

- 'static
- macro\_rules
- raw
- safe
- union

r[lex.keywords.weak.macro\_rules]

• macro\_rules is used to create custom macros.

r[lex.keywords.weak.union]

• union is used to declare a <u>union</u> and is only a keyword when used in a union declaration.

r[lex.keywords.weak.lifetime-static]

• 'static is used for the static lifetime and cannot be used as a <u>generic lifetime parameter</u> or <u>loop label</u>

```
// error[E0262]: invalid lifetime parameter name: `'static`
fn invalid_lifetime_parameter<'static>(s: &'static str) ->
&'static str { s }
```

r[lex.keywords.weak.safe]

• safe is used for functions and statics, which has meaning in <u>external</u> <u>blocks</u>.

r[lex.keywords.weak.raw]

 raw is used for <u>raw borrow operators</u>, and is only a keyword when matching a raw borrow operator form (such as &raw const expr or &raw mut expr). r[lex.keywords.weak.dyn.edition2018]

[!EDITION-2018] In the 2015 edition, <u>dyn</u> is a keyword when used in a type position followed by a path that does not start with **::** or **<**, a lifetime, a question mark, a <u>for</u> keyword or an opening parenthesis.

Beginning in the 2018 edition, dyn has been promoted to a strict keyword.

r[ident]

## **Identifiers**

XID\_Start -> <`XID\_Start` defined by Unicode>

XID\_Continue -> <`XID\_Continue` defined by Unicode>

RAW\_IDENTIFIER -> `r#` IDENTIFIER\_OR\_KEYWORD \_except `crate`, `self`, `super`, `Self`\_

NON\_KEYWORD\_IDENTIFIER -> IDENTIFIER\_OR\_KEYWORD \_except a
[strict][lex.keywords.strict] or [reserved]
[lex.keywords.reserved] keyword\_

IDENTIFIER -> NON\_KEYWORD\_IDENTIFIER | RAW\_IDENTIFIER

RESERVED\_RAW\_IDENTIFIER -> `r#\_`

r[ident.unicode] Identifiers follow the specification in <u>Unicode Standard</u> <u>Annex #31</u> for Unicode version 16.0, with the additions described below. Some examples of identifiers:

- foo
- \_identifier
- r#true
- Москва
- •

r[ident.profile] The profile used from UAX #31 is:

• Start := <u>XID Start</u>, plus the underscore character (U+005F)

- Continue := <u>XID Continue</u>
- Medial := empty

with the additional constraint that a single underscore character is not an identifier.

[!NOTE] Identifiers starting with an underscore are typically used to indicate an identifier that is intentionally unused, and will silence the unused warning in **rustc**.

r[ident.keyword] Identifiers may not be a <u>strict</u> or <u>reserved</u> keyword without the <u>r</u># prefix described below in <u>raw identifiers</u>.

r[ident.zero-width-chars] Zero width non-joiner (ZWNJ U+200C) and zero width joiner (ZWJ U+200D) characters are not allowed in identifiers.

r[ident.ascii-limitations] Identifiers are restricted to the ASCII subset of <u>XID Start</u> and <u>XID Continue</u> in the following situations:

- <u>extern crate</u> declarations (except the [AsClause] identifier)
- External crate names referenced in a <u>path</u>
- <u>Module</u> names loaded from the filesystem without a <u>path</u> attribute
- <u>no mangle</u> attributed items
- Item names in <u>external blocks</u>

r[ident.normalization]

### Normalization

Identifiers are normalized using Normalization Form C (NFC) as defined in <u>Unicode Standard Annex #15</u>. Two identifiers are equal if their NFC forms are equal.

<u>Procedural</u> and <u>declarative</u> macros receive normalized identifiers in their input.

r[ident.raw]

### **Raw identifiers**

r[ident.raw.intro] A raw identifier is like a normal identifier, but prefixed by **r**#. (Note that the **r**# prefix is not included as part of the actual identifier.)

r[ident.raw.allowed] Unlike a normal identifier, a raw identifier may be any strict or reserved keyword except the ones listed above for RAW\_IDENTIFIER.

r[ident.raw.reserved] It is an error to use the [RESERVED\_RAW\_IDENTIFIER] token r#\_ in order to avoid confusion with the [WildcardPattern].

r[comments]

## Comments

r[comments.syntax] @root LINE COMMENT -> `//` (~[`/` `!` LF] | `//`) ~LF\* | `//` BLOCK COMMENT -> `/\*` ( ~[`\*``!`] | `\*\*` | BlockCommentOrDoc ) ( BlockCommentOrDoc | ~`\*/` )\* `\*/` | `/\*\*/` | `/\*\*\*/` @root INNER\_LINE\_DOC -> `//!` ~[LF CR]\* INNER\_BLOCK\_DOC -> `/\*!` ( BlockCommentOrDoc | ~[`\*/` CR] )\* `\*/` @root OUTER LINE DOC -> `///` (~`/` ~[LF CR]\*)? OUTER\_BLOCK\_DOC -> `/\*\*` ( ~`\*` | BlockCommentOrDoc ) ( BlockCommentOrDoc | ~[`\*/` CR] )\* `\*/` @root BlockCommentOrDoc -> **BLOCK\_COMMENT** | OUTER BLOCK DOC | INNER\_BLOCK\_DOC

r[comments.normal]
## **Non-doc comments**

Comments follow the general C++ style of line (//) and block (/\* ... \*/) comment forms. Nested block comments are supported.

r[comments.normal.tokenization] Non-doc comments are interpreted as a form of whitespace.

r[comments.doc]

#### **Doc comments**

r[comments.doc.syntax] Line doc comments beginning with exactly *three* slashes (///), and block doc comments (/\*\* ... \*/), both outer doc comments, are interpreted as a special syntax for <u>doc attributes</u>.

r[comments.doc.attributes] That is, they are equivalent to writing # [doc="..."] around the body of the comment, i.e., /// Foo turns into # [doc="Foo"] and /\*\* Bar \*/ turns into #[doc="Bar"]. They must therefore appear before something that accepts an outer attribute.

r[comments.doc.inner-syntax] Line comments beginning with //! and block comments /\*! ... \*/ are doc comments that apply to the parent of the comment, rather than the item that follows.

r[comments.doc.inner-attributes] That is, they are equivalent to writing #![doc="..."] around the body of the comment. //! comments are usually used to document modules that occupy a source file.

r[comments.doc.bare-crs] The character U+000D (CR) is not allowed in doc comments.

[!NOTE] It is conventional for doc comments to contain Markdown, as expected by rustdoc. However, the comment syntax does not respect any internal Markdown. /\*\* `glob = "\*/\*.rs";` \*/ terminates the comment at the first \*/, and the remaining code would cause a syntax error. This slightly limits the content of block doc comments compared to line doc comments.

[!NOTE] The sequence U+000D (CR) immediately followed by U+000A (LF) would have been previously transformed into a single U+000A (LF).

### Examples

//! A doc comment that applies to the implicit anonymous
module of this crate

pub mod outer\_module {

//! - Inner line doc
//!! - Still an inner line doc (but with a bang at the
beginning)

/\*! - Inner block doc \*/
 /\*!! - Still an inner block doc (but with a bang at the
beginning) \*/

// - Only a comment
/// - Outer line doc (exactly 3 slashes)
//// - Only a comment
/\* - Only a comment \*/
/\*\* - Outer block doc (exactly) 2 asterisks \*/
/\*\*\* - Only a comment \*/
pub mod inner\_module {}

```
pub mod nested_comments {
    /* In Rust /* we can /* nest comments */ */ */
```

// All three types of block comments can contain or be
nested inside

// any other type:

/\* /\* \*/ /\*\* \*/ /\*! \*/ \*/
/\*! /\* \*/ /\*\* \*/ /\*! \*/ \*/
/\*

```
pub mod dummy_item {}
    }
    pub mod degenerate_cases {
        // empty inner line doc
        //!
        // empty inner block doc
        /*!*/
        // empty line comment
        11
        // empty outer line doc
        111
        // empty block comment
        /**/
        pub mod dummy_item {}
         // empty 2-asterisk block isn't a doc block, it is a
block comment
        /***/
    }
    /* The next one isn't allowed because outer doc comments
       require an item that will receive the doc */
    /// Where is my item?
    mod boo {}
#
}
```

r[lex.whitespace]

# Whitespace

r[lex.whitespace.intro] Whitespace is any non-empty string containing only characters that have the <u>Pattern White Space</u> Unicode property, namely:

- U+0009 (horizontal tab, '\t')
- U+000A (line feed, '\n')
- U+000B (vertical tab)
- U+000C (form feed)
- U+000D (carriage return, '\r')
- U+0020 (space, ' ')
- U+0085 (next line)
- U+200E (left-to-right mark)
- U+200F (right-to-left mark)
- U+2028 (line separator)
- U+2029 (paragraph separator)

r[lex.whitespace.token-sep] Rust is a "free-form" language, meaning that all forms of whitespace serve only to separate *tokens* in the grammar, and have no semantic significance.

r[lex.whitespace.replacement] A Rust program has identical meaning if each whitespace element is replaced with any other legal whitespace element, such as a single space character. r[lex.token]

# Tokens

r[lex.token.syntax]

Token ->

IDENTIFIER\_OR\_KEYWORD

- | RAW\_IDENTIFIER
- | CHAR\_LITERAL
- | STRING\_LITERAL
- | RAW\_STRING\_LITERAL
- | BYTE\_LITERAL
- | BYTE\_STRING\_LITERAL
- | RAW\_BYTE\_STRING\_LITERAL

| C\_STRING\_LITERAL

- | RAW\_C\_STRING\_LITERAL
- | INTEGER\_LITERAL
- | FLOAT\_LITERAL
- | LIFETIME\_TOKEN
- | PUNCTUATION
- | RESERVED\_TOKEN

r[lex.token.intro] Tokens are primitive productions in the grammar defined by regular (non-recursive) languages. Rust source input can be broken down into the following kinds of tokens:

- <u>Keywords</u>
- <u>Identifiers</u>
- <u>Literals</u>
- <u>Lifetimes</u>
- <u>Punctuation</u>
- <u>Delimiters</u>

Within this documentation's grammar, "simple" tokens are given in <u>string table production</u> form, and appear in monospace font.

r[lex.token.literal]

# Literals

Literals are tokens used in <u>literal expressions</u>.

# Examples

## **Characters and strings**

	Example	# set s <sup>1</sup>	Characters	Escapes
<u>Character</u>	'Η'	0	All Unicode	Quote&ASCII&Unicode
<u>String</u>	"hello "	0	All Unicode	Quote&ASCII&Unicode
<u>Raw</u> string	r#"hel lo"#	<256	All Unicode	N/A
<u>Byte</u>	b'H'	0	All ASCII	<u>Quote</u> & <u>Byte</u>
<u>Byte</u> string	b"hell o"	0	All ASCII	<u>Quote</u> & <u>Byte</u>
<u>Raw byte</u> <u>string</u>	br#"he llo"#	<256	All ASCII	N/A
<u>C string</u>	c"hell o"	0	All Unicode	Quote&Byte&Unicode
<u>Raw C</u> <u>string</u>	cr#"he llo"#	<256	All Unicode	N/A

1

The number of *#*s on each side of the same literal must be equivalent.

[!NOTE] Character and string literal tokens never include the sequence of U+000D (CR) immediately followed by U+000A (LF): this pair would have been previously transformed into a single U+000A (LF).

## **ASCII escapes**

	Name
\x41	7-bit character code (exactly 2 digits, up to 0x7F)
١n	Newline
١r	Carriage return
١t	Tab
	Backslash
\0	Null

## **Byte escapes**

	Name
\x7F	8-bit character code (exactly 2 digits)
\n	Newline
١r	Carriage return
١t	Tab
	Backslash
\0	Null

## **Unicode escapes**

	Name
\u{7FFF}	24-bit Unicode character code (up to 6 digits)

#### **Quote escapes**

	Name
$\mathbf{\lambda}^{\mathbf{r}}$	Single quote
$\sum_{i=1}^{n}$	Double quote

#### Numbers

<u>Number literals<sup>2</sup></u>	Example	Exponentiation
Decimal integer	98_222	N/A
Hex integer	0xff	N/A
Octal integer	0077	N/A
Binary integer	0b1111_0000	N/A
Floating-point	123.0E+77	Optional

2

All number literals allow \_ as a visual separator: 1\_234.0E+18f64

r[lex.token.literal.suffix]

#### Suffixes

r[lex.token.literal.literal.suffix.intro] A suffix is a sequence of characters following the primary part of a literal (without intervening whitespace), of the same form as a non-raw identifier or keyword.

r[lex.token.literal.suffix.syntax]
SUFFIX -> IDENTIFIER\_OR\_KEYWORD

```
SUFFIX_NO_E -> SUFFIX _not beginning with `e` or `E`_
```

r[lex.token.literal.suffix.validity] Any kind of literal (string, integer, etc) with any suffix is valid as a token.

A literal token with any suffix can be passed to a macro without producing an error. The macro itself will decide how to interpret such a

token and whether to produce an error or not. In particular, the literal fragment specifier for by-example macros matches literal tokens with arbitrary suffixes.

```
macro_rules! blackhole { ($tt:tt) => () }
macro_rules! blackhole_lit { ($l:literal) => () }
blackhole!("string"suffix); // OK
blackhole_lit!(1suffix); // OK
```

r[lex.token.literal.suffix.parse] However, suffixes on literal tokens which are interpreted as literal expressions or patterns are restricted. Any suffixes are rejected on non-numeric literal tokens, and numeric literal tokens are accepted only with suffixes from the list below.

Integer	Floating- point
u8, i8, u16, i16, u32, i32, u64, i64, u128,	f32, f64
i128, usize, isize	

# **Character and string literals**

r[lex.token.literal.char]

## **Character literals**

UNICODE\_ESCAPE ->

`\u{` ( HEX\_DIGIT `\_`\* ){1..6} `}`

r[lex.token.literal.char.intro] A *character literal* is a single Unicode character enclosed within two U+0027 (single-quote) characters, with the exception of U+0027 itself, which must be *escaped* by a preceding U+005C character ( $\backslash$ ).

r[lex.token.literal.str]

### **String literals**

r[lex.token.literal.str.syntax]
STRING\_LITERAL ->

```
`"` (
   ~[`"` `\` CR]
   QUOTE_ESCAPE
   ASCII_ESCAPE
   UNICODE_ESCAPE
   STRING_CONTINUE
)* `"` SUFFIX?
```

```
STRING_CONTINUE -> `\` LF
```

r[lex.token.literal.str.intro] A *string literal* is a sequence of any Unicode characters enclosed within two U+0022 (double-quote) characters, with the exception of U+0022 itself, which must be *escaped* by a preceding U+005C character ( $\backslash$ ).

r[lex.token.literal.str.linefeed] Line-breaks, represented by the character U+000A (LF), are allowed in string literals. When an unescaped U+005C character ( $\land$ ) occurs immediately before a line break, the line break does not appear in the string represented by the token. See <u>String continuation</u> escapes for details. The character U+000D (CR) may not appear in a string literal other than as part of such a string continuation escape.

r[lex.token.literal.char-escape]

#### **Character escapes**

r[lex.token.literal.char-escape.intro] Some additional *escapes* are available in either character or non-raw string literals. An escape starts with a U+005C ( $\backslash$ ) and continues with one of the following forms:

r[lex.token.literal.char-escape.ascii]

A 7-bit code point escape starts with U+0078 (x) and is followed by exactly two hex digits with value up to 0x7F. It denotes the ASCII character with value equal to the provided hex value. Higher values are not permitted because it is ambiguous whether they mean Unicode code points or byte values.

r[lex.token.literal.char-escape.unicode]

- A 24-bit code point escape starts with U+0075 (u) and is followed by up to six *hex digits* surrounded by braces U+007B ({) and U+007D (}). It denotes the Unicode code point equal to the provided hex value.
   r[lex.token.literal.char-escape.whitespace]
- A whitespace escape is one of the characters U+006E (n), U+0072 (r), or U+0074 (t), denoting the Unicode values U+000A (LF), U+000D (CR) or U+0009 (HT) respectively.

r[lex.token.literal.char-escape.null]

• The *null escape* is the character U+0030 (0) and denotes the Unicode value U+0000 (NUL).

r[lex.token.literal.char-escape.slash]

• The *backslash escape* is the character U+005C (\) which must be escaped in order to denote itself.

r[lex.token.literal.str-raw]

## **Raw string literals**

r[lex.token.literal.str-raw.syntax]
RAW\_STRING\_LITERAL -> `r` RAW\_STRING\_CONTENT SUFFIX?

RAW\_STRING\_CONTENT ->

`"` ( ~CR )\*? `"`

| `#` RAW\_STRING\_CONTENT `#`

r[lex.token.literal.str-raw.intro] Raw string literals do not process any escapes. They start with the character U+0072 (r), followed by fewer than 256 of the character U+0023 (#) and a U+0022 (double-quote) character.

r[lex.token.literal.str-raw.body] The *raw string body* can contain any sequence of Unicode characters other than U+000D (CR). It is terminated only by another U+0022 (double-quote) character, followed by the same number of U+0023 (#) characters that preceded the opening U+0022 (double-quote) character.

r[lex.token.literal.str-raw.content] All Unicode characters contained in the raw string body represent themselves, the characters U+0022 (double-quote) (except when followed by at least as many U+0023 (#) characters as were used to start the raw string literal) or U+005C (\) do not have any special meaning.

Examples for string literals:

"foo"; r"foo";	//	foo
"\"foo\"";	//	"foo"
"foo #\"# bar";		
r##"foo #"# bar"##;	//	foo #"# bar
"\x52"; "R"; r"R";	//	R
"\\x52"; r"\x52";	11	\x52

## Byte and byte string literals

r[lex.token.byte]

#### **Byte literals**

```
r[lex.token.byte.syntax]
BYTE_LITERAL ->
    `b'` ( ASCII_FOR_CHAR | BYTE_ESCAPE ) `'` SUFFIX?
```

ASCII\_FOR\_CHAR ->

<any ASCII (i.e. 0x00 to 0x7F) except `'`, `\`, LF, CR, or TAB>

BYTE\_ESCAPE ->

`\x` HEX\_DIGIT HEX\_DIGIT

|`\n`|`\r`|`\t`|`\\`|`\0`|`\'`|`\"`

r[lex.token.byte.intro] A *byte literal* is a single ASCII character (in the U+0000 to U+007F range) or a single *escape* preceded by the characters U+0062 (b) and U+0027 (single-quote), and followed by the character U+0027. If the character U+0027 is present within the literal, it must be *escaped* by a preceding U+005C (\) character. It is equivalent to a u8 unsigned 8-bit integer *number literal*.

r[lex.token.str-byte]

#### **Byte string literals**

r[lex.token.str-byte.syntax]

```
BYTE_STRING_LITERAL ->
```

`b"` ( ASCII\_FOR\_STRING | BYTE\_ESCAPE | STRING\_CONTINUE )\* `"` SUFFIX?

#### ASCII\_FOR\_STRING ->

<any ASCII (i.e 0x00 to 0x7F) except `"`, `\`, or CR>

r[lex.token.str-byte.intro] A non-raw *byte string literal* is a sequence of ASCII characters and *escapes*, preceded by the characters U+0062 (b) and U+0022 (double-quote), and followed by the character U+0022. If the character U+0022 is present within the literal, it must be *escaped* by a preceding U+005C (\) character. Alternatively, a byte string literal can be a *raw byte string literal*, defined below.

r[lex.token.str-byte.linefeed] Line-breaks, represented by the character U+000A (LF), are allowed in byte string literals. When an unescaped U+005C character ( $\land$ ) occurs immediately before a line break, the line break does not appear in the string represented by the token. See <u>String</u> <u>continuation escapes</u> for details. The character U+000D (CR) may not

appear in a byte string literal other than as part of such a string continuation escape.

r[lex.token.str-byte.escape] Some additional *escapes* are available in either byte or non-raw byte string literals. An escape starts with a U+005C ( $\setminus$ ) and continues with one of the following forms:

r[lex.token.str-byte.escape-byte]

• A *byte escape* escape starts with U+0078 (x) and is followed by exactly two *hex digits*. It denotes the byte equal to the provided hex value.

r[lex.token.str-byte.escape-whitespace]

A whitespace escape is one of the characters U+006E (n), U+0072 (r), or U+0074 (t), denoting the bytes values 0x0A (ASCII LF), 0x0D (ASCII CR) or 0x09 (ASCII HT) respectively.

r[lex.token.str-byte.escape-null]

• The *null escape* is the character U+0030 (0) and denotes the byte value 0×00 (ASCII NUL).

r[lex.token.str-byte.escape-slash]

• The *backslash escape* is the character U+005C (\) which must be escaped in order to denote its ASCII encoding 0x5C.

r[lex.token.str-byte-raw]

## **Raw byte string literals**

```
r[lex.token.str-byte-raw.syntax]
RAW_BYTE_STRING_LITERAL ->
    `br` RAW_BYTE_STRING_CONTENT SUFFIX?
RAW_BYTE_STRING_CONTENT ->
```

```
`"` ASCII_FOR_RAW*? `"`
```

| `#` RAW\_BYTE\_STRING\_CONTENT `#`

ASCII\_FOR\_RAW ->

<any ASCII (i.e. 0x00 to 0x7F) except CR>

r[lex.token.str-byte-raw.intro] Raw byte string literals do not process any escapes. They start with the character U+0062 (b), followed by U+0072 (r), followed by fewer than 256 of the character U+0023 (#), and a U+0022 (double-quote) character.

r[lex.token.str-byte-raw.body] The *raw string body* can contain any sequence of ASCII characters other than U+000D (CR). It is terminated only by another U+0022 (double-quote) character, followed by the same number of U+0023 (#) characters that preceded the opening U+0022 (double-quote) character. A raw byte string literal can not contain any non-ASCII byte.

r[lex.token.literal.str-byte-raw.content] All characters contained in the raw string body represent their ASCII encoding, the characters U+0022 (double-quote) (except when followed by at least as many U+0023 (#) characters as were used to start the raw string literal) or U+005C (\) do not have any special meaning.

Examples for byte string literals:

b"foo"; br"foo";	// foo
b"\"foo\"";	// "foo"
b"foo #\"# bar";	
br##"foo #"# bar"##;	// foo #"# bar
b"\x52";	// R
b"\\x52"; br"\x52";	// \x52

# C string and raw C string literals

```
r[lex.token.str-c]
```

#### **C** string literals

```
r[lex.token.str-c.syntax]
C_STRING_LITERAL ->
    `c"` (
```

```
~[`"``\` CR NUL]
| BYTE_ESCAPE _except `\0` or `\x00`_
| UNICODE_ESCAPE _except `\u{0}`, `\u{00}`, ...,
`\u{000000}`_
| STRING_CONTINUE
)* `"` SUFFIX?
```

r[lex.token.str-c.intro] A *C* string literal is a sequence of Unicode characters and *escapes*, preceded by the characters U+0063 (c) and U+0022 (double-quote), and followed by the character U+0022. If the character U+0022 is present within the literal, it must be *escaped* by a preceding U+005C (\) character. Alternatively, a C string literal can be a *raw C string literal*, defined below.

r[lex.token.str-c.null] C strings are implicitly terminated by byte  $0 \times 00$ , so the C string literal c"" is equivalent to manually constructing a &CStr from the byte string literal b"\x00". Other than the implicit terminator, byte  $0 \times 00$  is not permitted within a C string.

r[lex.token.str-c.linefeed] Line-breaks, represented by the character U+000A (LF), are allowed in C string literals. When an unescaped U+005C character ( $\backslash$ ) occurs immediately before a line break, the line break does not appear in the string represented by the token. See <u>String continuation</u> escapes for details. The character U+000D (CR) may not appear in a C string literal other than as part of such a string continuation escape.

r[lex.token.str-c.escape] Some additional *escapes* are available in nonraw C string literals. An escape starts with a U+005C (\) and continues with one of the following forms:

r[lex.token.str-c.escape-byte]

• A *byte escape* escape starts with U+0078 (x) and is followed by exactly two *hex digits*. It denotes the byte equal to the provided hex value.

r[lex.token.str-c.escape-unicode]

A 24-bit code point escape starts with U+0075 (u) and is followed by up to six *hex digits* surrounded by braces U+007B ({) and U+007D (}). It denotes the Unicode code point equal to the provided hex value, encoded as UTF-8.

r[lex.token.str-c.escape-whitespace]

A whitespace escape is one of the characters U+006E (n), U+0072 (r), or U+0074 (t), denoting the bytes values 0x0A (ASCII LF), 0x0D (ASCII CR) or 0x09 (ASCII HT) respectively.

r[lex.token.str-c.escape-slash]

• The *backslash escape* is the character U+005C (\) which must be escaped in order to denote its ASCII encoding 0x5C.

r[lex.token.str-c.char-unicode] A C string represents bytes with no defined encoding, but a C string literal may contain Unicode characters above U+007F. Such characters will be replaced with the bytes of that character's UTF-8 representation.

The following C string literals are equivalent:

```
c"æ"; // LATIN SMALL LETTER AE (U+00E6)
c"\u{00E6}";
c"\xC3\xA6";
```

```
r[lex.token.str-c.edition2021]
```

[!EDITION-2021] C string literals are accepted in the 2021 edition or later. In earlier additions the token c"" is lexed as c "".

```
r[lex.token.str-c-raw]
```

#### **Raw C string literals**

RAW\_C\_STRING\_CONTENT ->

`"` ( ~[CR NUL] )\*? `"`

| `#` RAW\_C\_STRING\_CONTENT `#`

r[lex.token.str-c-raw.intro] Raw C string literals do not process any escapes. They start with the character U+0063 (c), followed by U+0072 (r), followed by fewer than 256 of the character U+0023 (#), and a U+0022 (double-quote) character.

r[lex.token.str-c-raw.body] The *raw C string body* can contain any sequence of Unicode characters other than U+0000 (NUL) and U+000D (CR). It is terminated only by another U+0022 (double-quote) character, followed by the same number of U+0023 (#) characters that preceded the opening U+0022 (double-quote) character.

r[lex.token.str-c-raw.content] All characters contained in the raw C string body represent themselves in UTF-8 encoding. The characters U+0022(double-quote) (except when followed by at least as many U+0023 (#) characters as were used to start the raw C string literal) or U+005C (\) do not have any special meaning.

r[lex.token.str-c-raw.edition2021]

[!EDITION-2021] Raw C string literals are accepted in the 2021 edition or later. In earlier additions the token cr"" is lexed as cr "", and cr#""# is lexed as cr #""# (which is non-grammatical).

## **Examples for C string and raw C string literals**

c"foo"; cr"foo";	//	foo
c"\"foo\"";	//	"foo"
C"TOO #\"# Dar";		
cr##"foo #"# bar"##;	//	foo #"# baı
c"\x52"; c"R"; cr"R";	//	R
c"\\x52"; cr"\x52";	//	\x52

r[lex.token.literal.num]

# **Number literals**

A *number literal* is either an *integer literal* or a *floating-point literal*. The grammar for recognizing the two kinds of literals is mixed.

r[lex.token.literal.int]

## **Integer literals**

```
r[lex.token.literal.int.syntax]
INTEGER_LITERAL ->
  ( DEC_LITERAL | BIN_LITERAL | OCT_LITERAL | HEX_LITERAL )
SUFFIX_NO_E?
DEC_LITERAL -> DEC_DIGIT (DEC_DIGIT|`_`)*
BIN_LITERAL -> `0b` (BIN_DIGIT|`_`)* BIN_DIGIT (BIN_DIGIT|`_`)*
OCT_LITERAL -> `0o` (OCT_DIGIT|`_`)* OCT_DIGIT (OCT_DIGIT|`_`)*
HEX_LITERAL -> `0x` (HEX_DIGIT|`_`)* HEX_DIGIT (HEX_DIGIT|`_`)*
BIN_DIGIT -> [`0`-`1`]
OCT_DIGIT -> [`0`-`7`]
DEC_DIGIT -> [`0`-`9`]
HEX_DIGIT -> [`0`-`9` `a`-`f` `A`-`F`]
```

r[lex.token.literal.int.kind] An integer literal has one of four forms: r[lex.token.literal.int.kind-dec]

- A *decimal literal* starts with a *decimal digit* and continues with any mixture of *decimal digits* and *underscores*.
   r[lex.token.literal.int.kind-hex]
- A *hex literal* starts with the character sequence U+0030 U+0078 (0x) and continues as any mixture (with at least one digit) of hex digits and underscores.

r[lex.token.literal.int.kind-oct]

An octal literal starts with the character sequence U+0030 U+006F
 (00) and continues as any mixture (with at least one digit) of octal digits and underscores.

r[lex.token.literal.int.kind-bin]

A *binary literal* starts with the character sequence U+0030 U+0062
 (0b) and continues as any mixture (with at least one digit) of binary digits and underscores.

r[lex.token.literal.int.restriction] Like any literal, an integer literal may be followed (immediately, without any spaces) by a suffix as described above. The suffix may not begin with e or E, as that would be interpreted as the exponent of a floating-point literal. See <u>Integer literal expressions</u> for the effect of these suffixes.

Examples of integer literals which are accepted as literal expressions:

```
# #![allow(overflowing_literals)]
123;
123i32;
123u32;
123_u32;
0xff;
0xff_u8;
0x01_f32; // integer 7986, not floating-point 1.0
0x01_e3; // integer 483, not floating-point 1000.0
0o70;
0o70_i16;
0b1111_1111_1001_0000;
0b1111_1111_1001_0000i64;
0b______1;
```

Ousize;

// These are too big for their type, but are accepted as literal expressions.

128\_i8; 256\_u8;

// This is an integer literal, accepted as a floating-point
literal expression.
5f32;

Note that -118, for example, is analyzed as two tokens: - followed by 118.

Examples of integer literals which are not accepted as literal expressions:

```
# #[cfg(false)] {
OinvalidSuffix;
123AFB43;
Ob010a;
OxAB_CD_EF_GH;
Ob1111_f32;
# }
```

r[lex.token.literal.int.tuple-field]

#### **Tuple index**

```
r[lex.token.literal.int.tuple-field.syntax]
TUPLE_INDEX -> INTEGER_LITERAL
```

r[lex.token.literal.int.tuple-field.intro] A tuple index is used to refer to the fields of <u>tuples</u>, <u>tuple structs</u>, and <u>tuple variants</u>.

r[lex.token.literal.int.tuple-field.eq] Tuple indices are compared with the literal token directly. Tuple indices start with 0 and each successive index increments the value by 1 as a decimal value. Thus, only decimal values will match, and the value must not have any extra 0 prefix characters.

```
let example = ("dog", "cat", "horse");
let dog = example.0;
let cat = example.1;
// The following examples are invalid.
let cat = example.01; // ERROR no field named `01`
let horse = example.0b10; // ERROR no field named `0b10`
```

[!NOTE] Tuple indices may include certain suffixes, but this is not intended to be valid, and may be removed in a future version. See <u>https://github.com/rust-lang/rust/issues/60210</u> for more information.

r[lex.token.literal.float]

### **Floating-point literals**

```
r[lex.token.literal.float.syntax]
FLOAT_LITERAL ->
    DEC_LITERAL `.` _not immediately followed by `.`, `_` or
an XID_Start character_
```

```
| DEC_LITERAL `.` DEC_LITERAL SUFFIX_NO_E?
```

```
| DEC_LITERAL (`.` DEC_LITERAL)? FLOAT_EXPONENT SUFFIX?
```

r[lex.token.literal.float.form] A *floating-point literal* has one of two forms:

- A *decimal literal* followed by a period character U+002E ( . ). This is optionally followed by another decimal literal, with an optional *exponent*.
- A single *decimal literal* followed by an *exponent*.

r[lex.token.literal.float.suffix] Like integer literals, a floating-point literal may be followed by a suffix, so long as the pre-suffix part does not end with U+002E ( . ). The suffix may not begin with e or E if the literal does not include an exponent. See <u>Floating-point literal expressions</u> for the effect of these suffixes.

Examples of floating-point literals which are accepted as literal expressions:

```
123.0f64;
0.1f64;
0.1f32;
12E+99_f64;
let x: f64 = 2.;
```

This last example is different because it is not possible to use the suffix syntax with a floating point literal end.token.ing in a period. 2.f64 would attempt to call a method named f64 on 2.

Note that -1.0, for example, is analyzed as two tokens: - followed by 1.0.

Examples of floating-point literals which are not accepted as literal expressions:

```
# #[cfg(false)] {
2.0f80;
2e5f80;
2e5e6;
2.0e5e6;
1.3e10u64;
# }
```

r[lex.token.literal.reserved]

#### **Reserved forms similar to number literals**

```
r[lex.token.literal.reserved.syntax]
RESERVED_NUMBER ->
BIN_LITERAL [`2`-`9`]
| OCT_LITERAL [`8`-`9`]
| ( BIN_LITERAL | OCT_LITERAL | HEX_LITERAL ) `.` _not
immediately followed by `.`, `_` or an XID_Start character_
| ( BIN_LITERAL | OCT_LITERAL ) (`e`|`E`)
| `0b` `_`* <end of input or not BIN_DIGIT>
| `0o` `_`* <end of input or not OCT_DIGIT>
| `0x` `_`* <end of input or not HEX_DIGIT>
```

| DEC\_LITERAL ( `.` DEC\_LITERAL )? (`e` | `E`) (`+` | `-`)? <end of input or not DEC\_DIGIT>

r[lex.token.literal.reserved.intro] The following lexical forms similar to number literals are *reserved forms*. Due to the possible ambiguity these raise, they are rejected by the tokenizer instead of being interpreted as separate tokens.

r[lex.token.literal.reserved.out-of-range]

- An unsuffixed binary or octal literal followed, without intervening whitespace, by a decimal digit out of the range for its radix. r[lex.token.literal.reserved.period]
- An unsuffixed binary, octal, or hexadecimal literal followed, without intervening whitespace, by a period character (with the same restrictions on what follows the period as for floating-point literals).
   r[lex.token.literal.reserved.exp]
- An unsuffixed binary or octal literal followed, without intervening whitespace, by the character e or E.

r[lex.token.literal.reserved.empty-with-radix]

- Input which begins with one of the radix prefixes but is not a valid binary, octal, or hexadecimal literal (because it contains no digits). r[lex.token.literal.reserved.empty-exp]
- Input which has the form of a floating-point literal with no digits in the exponent.

Examples of reserved forms:

```
0b0102; // this is not `0b010` followed by `2`
0o1279; // this is not `0o127` followed by `9`
0x80.0; // this is not `0x80` followed by `.` and `0`
0b101e; // this is not a suffixed literal, or `0b101` followed
by `e`
0b; // this is not an integer literal, or `0` followed by
```

`b` Ob\_; // this is not an integer literal, or `0` followed by `b\_` 2e; // this is not a floating-point literal, or `2` followed by `e` 2.0e; // this is not a floating-point literal, or `2.0` followed by `e` 2em; // this is not a suffixed literal, or `2` followed by `em` 2.0em; // this is not a suffixed literal, or `2.0` followed by `em`

r[lex.token.life]

## Lifetimes and loop labels

r[lex.token.life.syntax]

LIFETIME\_TOKEN ->

`'` IDENTIFIER\_OR\_KEYWORD \_not immediately followed by

| `'\_` \_not immediately followed by `'`\_

| RAW\_LIFETIME

LIFETIME\_OR\_LABEL ->

`'` NON\_KEYWORD\_IDENTIFIER \_not immediately followed by
`'`\_

| RAW\_LIFETIME

RAW\_LIFETIME ->

`'r#` IDENTIFIER\_OR\_KEYWORD \_except `crate`, `self`, `super`, `Self` and not immediately followed by `'`\_

```
RESERVED_RAW_LIFETIME -> `'r#_` _not immediately followed by
`'`_
```

r[lex.token.life.intro] Lifetime parameters and <u>loop labels</u> use LIFETIME\_OR\_LABEL tokens. Any LIFETIME\_TOKEN will be accepted by the lexer, and for example, can be used in macros.

r[lex.token.life.raw.intro] A raw lifetime is like a normal lifetime, but its identifier is prefixed by **r**#. (Note that the **r**# prefix is not included as part of the actual lifetime.)

r[lex.token.life.raw.allowed] Unlike a normal lifetime, a raw lifetime may be any strict or reserved keyword except the ones listed above for RAW\_LIFETIME.

r[lex.token.life.raw.reserved] It is an error to use the RESERVED\_RAW\_LIFETIME token 'r#\_ in order to avoid confusion with the <u>placeholder lifetime</u>.

r[lex.token.life.raw.edition2021]

[!EDITION-2021] Raw lifetimes are accepted in the 2021 edition or later. In earlier additions the token 'r#lt is lexed as 'r # lt.

r[lex.token.punct]

# Punctuation

r[lex.token.punct.syntax]

PUNCTUATION	->
`=`	
`<`	
`<=`	
`==`	
`!=`	
`>=`	
`>`	
`&&`	
`  `	
`!`	
`~`	
`+`	
`-`	
`*`	
`/`	
`%`	
`^`	
`&`	
`<<`	
`>>`	
`+=`	
-=	
*=	
/=	
`%=`	
`^=`	
&=	
=	
`<<=`	
`>>=`	

) (	@	`		
	•	`		
	•	•	`	
	•	•	•	`
	•	•	=	`
	,	`		
	;	`		
	:	`		
	:	:	`	
	- :	>	`	
``	<	-	`	
`:	=:	>	`	
` <i></i>	ŧ	`		
`;	\$	`		
	?	`		
`_	_	`		
`.	{	`		
	}	`		
	[	`		
	]	`		
Ì	(	`		
	)	`		

r[lex.token.punct.intro] Punctuation symbol tokens are listed here for completeness. Their individual usages and meanings are defined in the linked pages.

Symbol	Name	Usage
+	Plus	<u>Addition, Trait Bounds, Macro Kleene</u> <u>Matcher</u>
-	Minus	Subtraction, Negation
*	Star	<u>Multiplication, Dereference, Raw</u> <u>Pointers, Macro Kleene Matcher, Use</u> <u>wildcards</u>
/	Slash	<u>Division</u>

Symbol	Name	Usage
%	Percent	<u>Remainder</u>
Λ	Caret	Bitwise and Logical XOR
!	Not	<u>Bitwise and Logical NOT, Macro Calls,</u> <u>Inner Attributes, Never Type, Negative</u> <u>impls</u>
&	And	<u>Bitwise and Logical AND, Borrow,</u> <u>References, Reference patterns</u>
	Or	<u>Bitwise and Logical OR, Closures,</u> Patterns in <u>match, if let</u> , and <u>while let</u>
&&	AndAnd	<u>Lazy AND, Borrow, References, Reference patterns</u>
	OrOr	<u>Lazy OR</u> , <u>Closures</u>
<<	Shl	Shift Left, Nested Generics
>>	Shr	Shift Right, Nested Generics
+=	PlusEq	Addition assignment
-=	MinusEq	Subtraction assignment
*=	StarEq	Multiplication assignment
/=	SlashEq	Division assignment
%=	PercentEq	Remainder assignment
^=	CaretEq	Bitwise XOR assignment
&=	AndEq	Bitwise And assignment
=	OrEq	Bitwise Or assignment
<<=	ShlEq	Shift Left assignment
>>=	ShrEq	Shift Right assignment, Nested Generics
=	Eq	<u>Assignment</u> , <u>Attributes</u> , Various type definitions

Symbol	Name	Usage
==	EqEq	<u>Equal</u>
!=	Ne	Not Equal
>	Gt	Greater than, Generics, Paths
<	Lt	Less than, Generics, Paths
>=	Ge	Greater than or equal to, Generics
<=	Le	Less than or equal to
@	At	Subpattern binding
-	Underscore	<u>Wildcard patterns</u> , <u>Inferred types</u> , Unnamed items in <u>constants</u> , <u>extern</u> <u>crates</u> , <u>use declarations</u> , and <u>destructuring assignment</u>
	Dot	<u>Field access</u> , <u>Tuple index</u>
	DotDot	<u>Range, Struct expressions, Patterns,</u> <u>Range Patterns</u>
	DotDotDot	Variadic functions, Range patterns
=	DotDotEq	<u>Inclusive Range, Range patterns</u>
,	Comma	Various separators
;	Semi	Terminator for various items and statements, <u>Array types</u>
:	Colon	Various separators
::	PathSep	Path separator
->	RArrow	<u>Function return type</u> , <u>Closure return</u> <u>type</u> , <u>Function pointer type</u>
=>	FatArrow	Match arms, Macros
<-	LArrow	The left arrow symbol has been unused since before Rust 1.0, but it is still treated as a single token

Symbol	Name	Usage
#	Pound	<u>Attributes</u>
\$	Dollar	<u>Macros</u>
?	Question	<u>Question mark operator, Questionably</u> <u>sized, Macro Kleene Matcher</u>
~	Tilde	The tilde operator has been unused since before Rust 1.0, but its token may still be used

r[lex.token.delim]

# **Delimiters**

Bracket punctuation is used in various parts of the grammar. An open bracket must always be paired with a close bracket. Brackets and the tokens within them are referred to as "token trees" in <u>macros</u>. The three types of brackets are:

Bracket	Туре
{ }	Curly braces
[]	Square brackets
( )	Parentheses

r[lex.token.reserved]
#### **Reserved tokens**

r[lex.token.reserved.intro] Several token forms are reserved for future use. It is an error for the source input to match one of these forms.

r[lex.token.reserved.syntax]

RESERVED\_TOKEN ->

RESERVED\_GUARDED\_STRING\_LITERAL

| RESERVED\_NUMBER

| RESERVED\_POUNDS

| RESERVED\_RAW\_IDENTIFIER

| RESERVED\_RAW\_LIFETIME

| RESERVED\_TOKEN\_DOUBLE\_QUOTE

| RESERVED\_TOKEN\_LIFETIME

| RESERVED\_TOKEN\_POUND

| RESERVED\_TOKEN\_SINGLE\_QUOTE

r[lex.token.reserved-prefix]

#### **Reserved prefixes**

r[lex.token.reserved-prefix.syntax]

RESERVED\_TOKEN\_DOUBLE\_QUOTE ->

( IDENTIFIER\_OR\_KEYWORD \_except `b` or `c` or `r` or `br` or `cr`\_ | `\_` ) `"`

RESERVED\_TOKEN\_SINGLE\_QUOTE ->

( IDENTIFIER\_OR\_KEYWORD \_except `b`\_ | `\_` ) `'`

#### RESERVED\_TOKEN\_POUND ->

( IDENTIFIER\_OR\_KEYWORD \_except `r` or `br` or `cr`\_ | `\_` ) `#`

#### RESERVED\_TOKEN\_LIFETIME ->

`'` ( IDENTIFIER\_OR\_KEYWORD \_except `r`\_ | `\_` ) `#`

r[lex.token.reserved-prefix.intro] Some lexical forms known as *reserved prefixes* are reserved for future use.

r[lex.token.reserved-prefix.id] Source input which would otherwise be lexically interpreted as a non-raw identifier (or a keyword or \_\_) which is immediately followed by a #, ', or " character (without intervening whitespace) is identified as a reserved prefix.

r[lex.token.reserved-prefix.raw-token] Note that raw identifiers, raw string literals, and raw byte string literals may contain a # character but are not interpreted as containing a reserved prefix.

r[lex.token.reserved-prefix.strings] Similarly the r, b, br, c, and cr prefixes used in raw string literals, byte literals, byte string literals, raw byte string literals, C string literals, and raw C string literals are not interpreted as reserved prefixes.

r[lex.token.reserved-prefix.life] Source input which would otherwise be lexically interpreted as a non-raw lifetime (or a keyword or \_\_) which is immediately followed by a # character (without intervening whitespace) is identified as a reserved lifetime prefix. r[lex.token.reserved-prefix.edition2021]

[!EDITION-2021] Starting with the 2021 edition, reserved prefixes are reported as an error by the lexer (in particular, they cannot be passed to macros).

Before the 2021 edition, reserved prefixes are accepted by the lexer and interpreted as multiple tokens (for example, one token for the identifier or keyword, followed by a *#* token).

Examples accepted in all editions:

```
macro_rules! lexes {($($_:tt)*) => {}}
lexes!{a #foo}
lexes!{continue 'foo}
lexes!{match "..." {}}
lexes!{r#let#foo} // three tokens: r#let # foo
lexes!{'prefix #lt}
```

```
Examples accepted before the 2021 edition but rejected later:
macro_rules! lexes {($($_:tt)*) => {}}
lexes!{a#foo}
lexes!{continue'foo}
lexes!{match"..." {}}
lexes!{'prefix#lt}
```

r[lex.token.reserved-guards]

#### **Reserved guards**

r[lex.token.reserved-guards.syntax]
RESERVED\_GUARDED\_STRING\_LITERAL -> `#`+ STRING\_LITERAL

```
RESERVED_POUNDS -> `#`{2..}
```

r[lex.token.reserved-guards.intro] The reserved guards are syntax reserved for future use, and will generate a compile error if used.

r[lex.token.reserved-guards.string-literal] The *reserved guarded string literal* is a token of one or more U+0023 (#) immediately followed by a [STRING\_LITERAL].

r[lex.token.reserved-guards.pounds] The *reserved pounds* is a token of two or more U+0023 (#).

r[lex.token.reserved-guards.edition2024]

[!EDITION-2024] Before the 2024 edition, reserved guards are accepted by the lexer and interpreted as multiple tokens. For example, the **#"foo"#** form is interpreted as three tokens. **##** is interpreted as two tokens.

r[macro]

# Macros

r[macro.intro] The functionality and syntax of Rust can be extended with custom definitions called macros. They are given names, and invoked through a consistent syntax: some\_extension!(...).

There are two ways to define new macros:

- <u>Macros by Example</u> define new syntax in a higher-level, declarative way.
- <u>Procedural Macros</u> define function-like macros, custom derives, and custom attributes using functions that operate on input tokens.

r[macro.invocation]

## **Macro Invocation**

r[macro.invocation.intro] A macro invocation expands a macro at compile time and replaces the invocation with the result of the macro. Macros may be invoked in the following situations:

r[macro.invocation.expr]

• <u>Expressions</u> and <u>statements</u> r[macro.invocation.pattern]

- <u>Patterns</u> r[macro.invocation.type]
- <u>Types</u>

r[macro.invocation.item]

• <u>Items</u> including <u>associated items</u> r[macro.invocation.nested] • <u>macro rules</u> transcribers r[macro.invocation.extern]

#### • External blocks

r[macro.invocation.item-statement] When used as an item or a statement, the [MacroInvocationSemi] form is used where a semicolon is required at the end when not using curly braces. <u>Visibility qualifiers</u> are never allowed before a macro invocation or <u>macro rules</u> definition.

```
// Used as an expression.
let x = vec![1, 2, 3];
// Used as a statement.
println!("Hello!");
// Used in a pattern.
macro_rules! pat {
    ($i:ident) => (Some($i))
}
if let pat!(x) = Some(1) {
    assert_eq!(x, 1);
}
// Used in a type.
macro_rules! Tuple {
    { $A:ty, $B:ty } => { ($A, $B) };
}
type N2 = Tuple!(i32, i32);
// Used as an item.
# use std::cell::RefCell;
thread_local!(static F00: RefCell<u32> = RefCell::new(1));
// Used as an associated item.
```

```
macro_rules! const_maker {
   ($t:ty, $v:tt) => { const CONST: $t = $v; };
}
trait T {
   const_maker!{i32, 7}
}
// Macro calls within macros.
macro_rules! example {
   () => { println!("Macro call in a macro!") };
}
// Outer macro `example` is expanded, then inner macro
`println` is expanded.
example!();
```

r[macro.decl]

# **Macros By Example**

```
r[macro.decl.syntax]
MacroRulesDefinition ->
    `macro rules` `!` IDENTIFIER MacroRulesDef
MacroRulesDef ->
     `(` MacroRules `)` `;`
    | `[` MacroRules `]` `;`
    | `{` MacroRules `}`
MacroRules ->
   MacroRule ( `;` MacroRule )* `;`?
MacroRule ->
    MacroMatcher `=>` MacroTranscriber
MacroMatcher ->
     `(` MacroMatch* `)`
    | `[` MacroMatch* `]`
    | `{` MacroMatch* `}`
MacroMatch ->
      Token _except `$` and [delimiters][lex.token.delim]_
    | MacroMatcher
        | `$` ( IDENTIFIER_OR_KEYWORD _except `crate`_ |
RAW_IDENTIFIER | `_` ) `:` MacroFragSpec
    | `$` `(` MacroMatch+ `)` MacroRepSep? MacroRepOp
MacroFragSpec ->
        `block` | `expr` | `expr_2021` | `ident` | `item` |
`lifetime` | `literal`
    | `meta` | `pat` | `pat_param` | `path` | `stmt` | `tt` |
`ty` | `vis`
```

MacroRepSep -> Token \_except [delimiters][lex.token.delim] and
[MacroRepOp]\_

MacroRepOp -> `\*` | `+` | `?`

#### MacroTranscriber -> DelimTokenTree

r[macro.decl.intro] macro\_rules allows users to define syntax extension in a declarative way. We call such extensions "macros by example" or simply "macros".

Each macro by example has a name, and one or more *rules*. Each rule has two parts: a *matcher*, describing the syntax that it matches, and a *transcriber*, describing the syntax that will replace a successfully matched invocation. Both the matcher and the transcriber must be surrounded by delimiters. Macros can expand to expressions, statements, items (including traits, impls, and foreign items), types, or patterns.

r[macro.decl.transcription]

## Transcribing

r[macro.decl.transcription.intro] When a macro is invoked, the macro expander looks up macro invocations by name, and tries each macro rule in turn. It transcribes the first successful match; if this results in an error, then future matches are not tried.

r[macro.decl.transcription.lookahead] When matching, no lookahead is performed; if the compiler cannot unambiguously determine how to parse the macro invocation one token at a time, then it is an error. In the following example, the compiler does not look ahead past the identifier to see if the following token is a ), even though that would allow it to parse the invocation unambiguously:

```
macro_rules! ambiguity {
    ($($i:ident)* $j:ident) => { };
}
```

```
ambiguity!(error); // Error: local ambiguity
```

r[macro.decl.transcription.syntax] In both the matcher and the transcriber, the \$ token is used to invoke special behaviours from the macro engine (described below in <u>Metavariables</u> and <u>Repetitions</u>). Tokens that aren't part of such an invocation are matched and transcribed literally, with one exception. The exception is that the outer delimiters for the matcher will match any pair of delimiters. Thus, for instance, the matcher (()) will match {()} but not {{}. The character \$ cannot be matched or transcribed literally.

r[macro.decl.transcription.fragment]

## Forwarding a matched fragment

When forwarding a matched fragment to another macro-by-example, matchers in the second macro will see an opaque AST of the fragment type. The second macro can't use literal tokens to match the fragments in the matcher, only a fragment specifier of the same type. The ident, lifetime, and tt fragment types are an exception, and *can* be matched by literal tokens. The following illustrates this restriction:

```
foo!(3);
```

The following illustrates how tokens can be directly matched after matching a tt fragment:

```
// compiles OK
macro_rules! foo {
    ($1:tt) => { bar!($1); }
}
macro_rules! bar {
    (3) => {}
}
foo!(3);
```

r[macro.decl.meta]

## **Metavariables**

r[macro.decl.meta.intro] In the matcher, **\$** *name* : *fragment-specifier* matches a Rust syntax fragment of the kind specified and binds it to the metavariable **\$** *name*.

r[macro.decl.meta.specifier] Valid fragment specifiers are:

- block: a [BlockExpression]
- expr: an [Expression]
- expr\_2021: an [Expression] except [UnderscoreExpression] and [ConstBlockExpression] (see [macro.decl.meta.edition2024])
- ident: an [IDENTIFIER\_OR\_KEYWORD], [RAW\_IDENTIFIER], or <u>\$crate</u>
- item:an[Item]
- lifetime: a [LIFETIME\_TOKEN]
- literal: matches <sup>?</sup>[LiteralExpression]
- meta: an [Attr], the contents of an attribute
- pat : a [Pattern] (see [macro.decl.meta.edition2021])
- pat\_param: a [PatternNoTopAlt]
- path: a [TypePath] style path
- **stmt**: a [Statement][grammar-Statement] without the trailing semicolon (except for item statements that require semicolons)
- tt: a [TokenTree] (a single <u>token</u> or tokens in matching delimiters
   (), [], or {})
- ty: a [Type][grammar-Type]
- vis: a possibly empty [Visibility] qualifier

r[macro.decl.meta.transcription] In the transcriber, metavariables are referred to simply by *sname*, since the fragment kind is specified in the matcher. Metavariables are replaced with the syntax element that matched them. Metavariables can be transcribed more than once or not at all.

r[macro.decl.meta.dollar-crate] The keyword metavariable <u>\$crate</u> can be used to refer to the current crate.

r[macro.decl.meta.edition2021]

[!EDITION-2021] Starting with the 2021 edition, pat fragmentspecifiers match top-level or-patterns (that is, they accept [Pattern]).

Before the 2021 edition, they match exactly the same fragments as pat\_param (that is, they accept [PatternNoTopAlt]).

The relevant edition is the one in effect for the macro\_rules! definition.

r[macro.decl.meta.edition2024]

[!EDITION-2024] Before the 2024 edition, expr fragment specifiers do not match [UnderscoreExpression] or [ConstBlockExpression] at the top level. They are allowed within subexpressions.

The expr\_2021 fragment specifier exists to maintain backwards compatibility with editions before 2024.

r[macro.decl.repetition]

#### Repetitions

r[macro.decl.repetition.intro] In both the matcher and transcriber, repetitions are indicated by placing the tokens to be repeated inside (...), followed by a repetition operator, optionally with a separator token between.

r[macro.decl.repetition.separator] The separator token can be any token
other than a delimiter or one of the repetition operators, but ; and , are the
most common. For instance, \$( \$i:ident ),\* represents any number of
identifiers separated by commas. Nested repetitions are permitted.

r[macro.decl.repetition.operators] The repetition operators are:

- \* --- indicates any number of repetitions.
- + --- indicates any number but at least one.
- ? --- indicates an optional fragment with zero or one occurrence.

r[macro.decl.repetition.optional-restriction] Since ? represents at most one occurrence, it cannot be used with a separator.

r[macro.decl.repetition.fragment] The repeated fragment both matches and transcribes to the specified number of the fragment, separated by the separator token. Metavariables are matched to every repetition of their corresponding fragment. For instance, the \$( \$i:ident ),\* example above matches \$i to all of the identifiers in the list.

During transcription, additional restrictions apply to repetitions so that the compiler knows how to expand them properly:

- 1. A metavariable must appear in exactly the same number, kind, and nesting order of repetitions in the transcriber as it did in the matcher. So for the matcher \$( \$i:ident ), \*, the transcribers => { \$i }, => { \$( \$( \$i)\* )\* }, and => { \$( \$i )+ } are all illegal, but => { \$( \$i );\* } is correct and replaces a comma-separated list of identifiers with a semicolon-separated list.
- 2. Each repetition in the transcriber must contain at least one metavariable to decide how many times to expand it. If multiple

metavariables appear in the same repetition, they must be bound to the same number of fragments. For instance, (\$(\$i:ident), \*; \$(\$j:ident), \*) => ((\$((\$i,\$j)), \*)) must bind the same number of \$i fragments as \$j fragments. This means that invoking the macro with (a, b, c; d, e, f) is legal and expands to ((a,d), (b,e), (c,f)), but (a, b, c; d, e) is illegal because it does not have the same number. This requirement applies to every layer of nested repetitions.

r[macro.decl.scope]

## Scoping, Exporting, and Importing

r[macro.decl.scope.intro] For historical reasons, the scoping of macros by example does not work entirely like items. Macros have two forms of scope: textual scope, and path-based scope. Textual scope is based on the order that things appear in source files, or even across multiple files, and is the default scoping. It is explained further below. Path-based scope works exactly the same way that item scoping does. The scoping, exporting, and importing of macros is controlled largely by attributes.

r[macro.decl.scope.unqualified] When a macro is invoked by an unqualified identifier (not part of a multi-part path), it is first looked up in textual scoping. If this does not yield any results, then it is looked up in path-based scoping. If the macro's name is qualified with a path, then it is only looked up in path-based scoping.

```
use lazy_static::lazy_static; // Path-based import.
macro_rules! lazy_static { // Textual definition.
        (lazy) => {};
}
```

```
lazy_static!{lazy} // Textual lookup finds our macro first.
self::lazy_static!{} // Path-based lookup ignores our macro,
finds imported one.
```

r[macro.decl.scope.textual]

# **Textual Scope**

r[macro.decl.scope.textual.intro] Textual scope is based largely on the order that things appear in source files, and works similarly to the scope of local variables declared with let except it also applies at the module level. When macro\_rules! is used to define a macro, the macro enters the scope after the definition (note that it can still be used recursively, since names are looked up from the invocation site), up until its surrounding scope, typically a module, is closed. This can enter child modules and even span across multiple files:

```
//// src/lib.rs
mod has_macro {
    // m!{} // Error: m is not in scope.
    macro_rules! m {
        () => {};
    }
    m!{} // OK: appears after declaration of m.
    mod uses_macro;
}
/// m!{} // Error: m is not in scope.
//// src/has_macro/uses_macro.rs
```

m!{} // OK: appears after declaration of m in src/lib.rs

r[macro.decl.scope.textual.shadow] It is not an error to define a macro multiple times; the most recent declaration will shadow the previous one unless it has gone out of scope.

```
macro_rules! m {
   (1) => {};
}
m!(1);
mod inner {
   m!(1);
   macro_rules! m {
      (2) => {};
   }
   // m!(1); // Error: no rule matches '1'
   m!(2);
```

```
macro_rules! m {
    (3) => {};
    }
    m!(3);
}
m!(1);
```

Macros can be declared and used locally inside functions as well, and work similarly:

```
fn foo() {
    // m!(); // Error: m is not in scope.
    macro_rules! m {
        () => {};
    }
    m!();
}
// m!(); // Error: m is not in scope.
```

```
r[macro.decl.scope.macro_use]
```

#### The macro\_use attribute

r[macro.decl.scope.macro\_use.mod-decl] The *macro\_use* attribute has two purposes. First, it can be used to make a module's macro scope not end when the module is closed, by applying it to a module:

```
#[macro_use]
mod inner {
    macro_rules! m {
        () => {};
    }
}
m!();
```

r[macro.decl.scope.macro\_use.prelude] Second, it can be used to import macros from another crate, by attaching it to an extern crate declaration

appearing in the crate's root module. Macros imported this way are imported into the <u>macro use prelude</u>, not textually, which means that they can be shadowed by any other name. While macros imported by # [macro\_use] can be used before the import statement, in case of a conflict, the last macro imported wins. Optionally, a list of macros to import can be specified using the [MetaListIdents] syntax; this is not supported when # [macro\_use] is applied to a module.

```
#[macro_use(lazy_static)] // Or #[macro_use] to import all
macros.
```

```
extern crate lazy_static;
```

```
lazy_static!{}
```

// self::lazy\_static!{} // Error: lazy\_static is not defined in
`self`

r[macro.decl.scope.macro\_use.export] Macros to be imported with #
[macro\_use] must be exported with #[macro\_export], which is described
below.

r[macro.decl.scope.path]

## **Path-Based Scope**

r[macro.decl.scope.path.intro] By default, a macro has no path-based scope. However, if it has the #[macro\_export] attribute, then it is declared in the crate root scope and can be referred to normally as such:

```
self::m!();
m!(); // OK: Path-based lookup finds m in the current module.
mod inner {
    super::m!();
    crate::m!();
}
mod mac {
    #[macro_export]
    macro_rules! m {
```

```
() => {};
}
}
```

r[macro.decl.scope.path.export] Macros labeled with #[macro\_export]
are always pub and can be referred to by other crates, either by path or by
#[macro\_use] as described above.

r[macro.decl.hygiene]

# Hygiene

r[macro.decl.hygiene.intro] Macros by example have *mixed-site hygiene*. This means that <u>loop labels</u>, <u>block labels</u>, and local variables are looked up at the macro definition site while other symbols are looked up at the macro invocation site. For example:

```
let x = 1;
fn func() {
    unreachable!("this is never called")
}
macro_rules! check {
    () => {
          assert_eq!(x, 1); // Uses `x` from the definition
site.
        func();
                           // Uses `func` from the invocation
site.
    };
}
{
    let x = 2;
    fn func() { /* does not panic */ }
    check!();
}
```

Labels and local variables defined in macro expansion are not shared between invocations, so this code doesn't compile:

```
macro_rules! m {
    (define) => {
        let x = 1;
    };
    (refer) => {
        dbg!(x);
    };
}
```

```
m!(define);
m!(refer);
```

r[macro.decl.hygiene.crate] A special case is the **\$crate** metavariable. It refers to the crate defining the macro, and can be used at the start of the path to look up items or macros which are not in scope at the invocation site.

```
//// Definitions in the `helper_macro` crate.
#[macro export]
macro_rules! helped {
    // () => { helper!() } // This might lead to an error due
to 'helper' not being in scope.
    () => { $crate::helper!() }
}
#[macro export]
macro_rules! helper {
    () => { () }
}
//// Usage in another crate.
// Note that `helper_macro::helper` is not imported!
use helper macro::helped;
fn unit() {
    helped!();
}
```

Note that, because **\$crate** refers to the current crate, it must be used with a fully qualified module path when referring to non-macro items:

```
pub mod inner {
    #[macro_export]
    macro_rules! call_foo {
        () => { $crate::inner::foo() };
    }
```

```
pub fn foo() {}
}
```

r[macro.decl.hygiene.vis] Additionally, even though \$crate allows a macro to refer to items within its own crate when expanding, its use has no effect on visibility. An item or macro referred to must still be visible from the invocation site. In the following example, any attempt to invoke call\_foo!() from outside its crate will fail because foo() is not public.

```
#[macro_export]
macro_rules! call_foo {
   () => { $crate::foo() };
}
```

```
fn foo() {}
```

**Version differences**: Prior to Rust 1.30, **\$crate** and **local\_inner\_macros** (below) were unsupported. They were added alongside path-based imports of macros (described above), to ensure that helper macros did not need to be manually imported by users of a macro-exporting crate. Crates written for earlier versions of Rust that use helper macros need to be modified to use **\$crate** or **local\_inner\_macros** to work well with path-based imports.

r[macro.decl.hygiene.local\_inner\_macros] When a macro is exported, the #[macro\_export] attribute can have the local\_inner\_macros keyword added to automatically prefix all contained macro invocations with \$crate::. This is intended primarily as a tool to migrate code written before \$crate was added to the language to work with Rust 2018's pathbased imports of macros. Its use is discouraged in new code.

```
#[macro_export(local_inner_macros)]
macro_rules! helped {
    () => { helper!() } // Automatically converted to
$crate::helper!().
}
```

```
#[macro_export]
macro_rules! helper {
   () => { () }
}
```

r[macro.decl.follow-set]

### **Follow-set Ambiguity Restrictions**

r[macro.decl.follow-set.intro] The parser used by the macro system is reasonably powerful, but it is limited in order to prevent ambiguity in current or future versions of the language.

r[macro.decl.follow-set.token-restriction] In particular, in addition to the rule about ambiguous expansions, a nonterminal matched by a metavariable must be followed by a token which has been decided can be safely used after that kind of match.

As an example, a macro matcher like *\$i:expr[,]* could in theory be accepted in Rust today, since [,] cannot be part of a legal expression and therefore the parse would always be unambiguous. However, because [ can start trailing expressions, [ is not a character which can safely be ruled out as coming after an expression. If [,] were accepted in a later version of Rust, this matcher would become ambiguous or would misparse, breaking working code. Matchers like *\$i:expr,* or *\$i:expr;* would be legal, however, because , and ; are legal expression separators. The specific rules are:

r[macro.decl.follow-set.token-expr-stmt]

- expr and stmt may only be followed by one of: =>, , , or ; .
   r[macro.decl.follow-set.token-pat\_param]
- pat\_param may only be followed by one of: =>, , , =, |, if, or in.
   r[macro.decl.follow-set.token-pat]
- pat may only be followed by one of: =>, , , =, if, or in.
   r[macro.decl.follow-set.token-path-ty]
- path and ty may only be followed by one of: =>, , , =, |, ;, :, >,
   >>, [, {, as, where, or a macro variable of block fragment specifier.

```
r[macro.decl.follow-set.token-vis]
```

vis may only be followed by one of: , , an identifier other than a non-raw priv, any token that can begin a type, or a metavariable with a ident, ty, or path fragment specifier.

r[macro.decl.follow-set.token-other]

• All other fragment specifiers have no restrictions. r[macro.decl.follow-set.edition2021]

[!EDITION-2021] Before the 2021 edition, pat may also be followed by [.

r[macro.decl.follow-set.repetition] When repetitions are involved, then the rules apply to every possible number of expansions, taking separators into account. This means:

- If the repetition includes a separator, that separator must be able to follow the contents of the repetition.
- If the repetition can repeat multiple times (\* or +), then the contents must be able to follow themselves.
- The contents of the repetition must be able to follow whatever comes before, and whatever comes after must be able to follow the contents of the repetition.
- If the repetition can match zero times (\* or ?), then whatever comes after must be able to follow whatever comes before.

For more detail, see the <u>formal specification</u>.

r[macro.proc]

# **Procedural Macros**

r[macro.proc.intro] *Procedural macros* allow creating syntax extensions as execution of a function. Procedural macros come in one of three flavors:

- <u>Function-like macros</u> custom!(...)
- <u>Derive macros</u> #[derive(CustomDerive)]
- <u>Attribute macros</u> #[CustomAttribute]

Procedural macros allow you to run code at compile time that operates over Rust syntax, both consuming and producing Rust syntax. You can sort of think of procedural macros as functions from an AST to another AST.

r[macro.proc.def] Procedural macros must be defined in the root of a crate with the <u>crate type</u> of proc-macro. The macros may not be used from the crate where they are defined, and can only be used when imported in another crate.

[!NOTE] When using Cargo, Procedural macro crates are defined with the proc-macro key in your manifest:

[lib] proc-macro = true

r[macro.proc.result] As functions, they must either return syntax, panic, or loop endlessly. Returned syntax either replaces or adds the syntax depending on the kind of procedural macro. Panics are caught by the compiler and are turned into a compiler error. Endless loops are not caught by the compiler which hangs the compiler.

Procedural macros run during compilation, and thus have the same resources that the compiler has. For example, standard input, error, and output are the same that the compiler has access to. Similarly, file access is the same. Because of this, procedural macros have the same security concerns that <u>Cargo's build scripts</u> have.

r[macro.proc.error] Procedural macros have two ways of reporting errors. The first is to panic. The second is to emit a [compile\_error] macro invocation. r[macro.proc.proc\_macro]

#### The proc\_macro crate

r[macro.proc\_macro.intro] Procedural macro crates almost always will link to the compiler-provided <u>proc\_macro\_crate</u>. The <u>proc\_macro</u> crate provides types required for writing procedural macros and facilities to make it easier.

r[macro.proc\_proc\_macro.token-stream] This crate primarily contains a <u>TokenStream</u> type. Procedural macros operate over *token streams* instead of AST nodes, which is a far more stable interface over time for both the compiler and for procedural macros to target. A *token stream* is roughly equivalent to Vec<TokenTree> where a TokenTree can roughly be thought of as lexical token. For example foo is an Ident token, . is a Punct token, and 1.2 is a Literal token. The TokenStream type, unlike Vec<TokenTree>, is cheap to clone.

r[macro.proc\_proc\_macro.span] All tokens have an associated Span. A Span is an opaque value that cannot be modified but can be manufactured. Spans represent an extent of source code within a program and are primarily used for error reporting. While you cannot modify a Span itself, you can always change the Span *associated* with any token, such as through getting a Span from another token.

r[macro.proc.hygiene]

#### **Procedural macro hygiene**

Procedural macros are *unhygienic*. This means they behave as if the output token stream was simply written inline to the code it's next to. This means that it's affected by external items and also affects external imports.

Macro authors need to be careful to ensure their macros work in as many contexts as possible given this limitation. This often includes using absolute paths to items in libraries (for example, ::std::option::Option instead of Option) or by ensuring that generated functions have names that are unlikely to clash with other functions (like \_\_internal\_foo instead of foo).

r[macro.proc.function]

### **Function-like procedural macros**

r[macro.proc.function.intro] *Function-like procedural macros* are procedural macros that are invoked using the macro invocation operator (!).

r[macro.proc.function.def] These macros are defined by a <u>public</u> <u>function</u> with the proc\_macro <u>attribute</u> and a signature of (TokenStream) -> TokenStream. The input <u>TokenStream</u> is what is inside the delimiters of the macro invocation and the output <u>TokenStream</u> replaces the entire macro invocation.

r[macro.proc.function.namespace] The proc\_macro attribute defines the macro in the macro namespace in the root of the crate.

For example, the following macro definition ignores its input and outputs a function answer into its scope.

```
# #![crate_type = "proc-macro"]
extern crate proc_macro;
use proc_macro::TokenStream;
#[proc_macro]
pub fn make_answer(_item: TokenStream) -> TokenStream {
    "fn answer() -> u32 { 42 }".parse().unwrap()
}
```

And then we use it in a binary crate to print "42" to standard output. extern crate proc\_macro\_examples; use proc\_macro\_examples::make\_answer;

```
make_answer!();
fn main() {
    println!("{}", answer());
}
```

r[macro.proc.function.invocation] Function-like procedural macros may be invoked in any macro invocation position, which includes <u>statements</u>,

<u>expressions</u>, <u>patterns</u>, <u>type expressions</u>, <u>item</u> positions, including items in <u>extern blocks</u>, inherent and trait <u>implementations</u>, and <u>trait definitions</u>.

r[macro.proc.derive]
#### **Derive macros**

r[macro.proc.derive.intro] *Derive macros* define new inputs for the <u>derive attribute</u>. These macros can create new <u>items</u> given the token stream of a <u>struct</u>, <u>enum</u>, or <u>union</u>. They can also define <u>derive macro helper</u> <u>attributes</u>.

r[macro.proc.derive.def] Custom derive macros are defined by a <u>public</u> <u>function</u> with the proc\_macro\_derive attribute and a signature of (TokenStream) -> TokenStream.

r[macro.proc.derive.namespace] The proc\_macro\_derive attribute defines the custom derive in the macro namespace in the root of the crate.

r[macro.proc.derive.output] The input <u>TokenStream</u> is the token stream of the item that has the <u>derive</u> attribute on it. The output <u>TokenStream</u> must be a set of items that are then appended to the <u>module</u> or <u>block</u> that the item from the input <u>TokenStream</u> is in.

The following is an example of a derive macro. Instead of doing anything useful with its input, it just appends a function answer.

```
# #![crate_type = "proc-macro"]
extern crate proc_macro;
use proc_macro::TokenStream;
#[proc_macro_derive(AnswerFn)]
pub fn derive_answer_fn(_item: TokenStream) -> TokenStream {
    "fn answer() -> u32 { 42 }".parse().unwrap()
}
And then using said derive macro:
```

extern crate proc\_macro\_examples; use proc\_macro\_examples::AnswerFn;

```
#[derive(AnswerFn)]
struct Struct;
```

```
fn main() {
```

```
assert_eq!(42, answer());
}
```

```
r[macro.proc.derive.attributes]
```

# Derive macro helper attributes

r[macro.proc.derive.attributes.intro] Derive macros can add additional <u>attributes</u> into the scope of the <u>item</u> they are on. Said attributes are called *derive macro helper attributes*. These attributes are <u>inert</u>, and their only purpose is to be fed into the derive macro that defined them. That said, they can be seen by all macros.

r[macro.proc.derive.attributes.def] The way to define helper attributes is to put an attributes key in the proc\_macro\_derive macro with a comma separated list of identifiers that are the names of the helper attributes.

For example, the following derive macro defines a helper attribute helper, but ultimately doesn't do anything with it.

```
# #![crate_type="proc-macro"]
# extern crate proc_macro;
# use proc_macro::TokenStream;
#[proc_macro_derive(HelperAttr, attributes(helper))]
pub fn derive_helper_attr(_item: TokenStream) -> TokenStream {
    TokenStream::new()
}
And then usage on the derive macro on a struct:
#[derive(HelperAttr)]
struct Struct {
```

```
#[helper] field: ()
```

```
}
```

```
r[macro.proc.attribute]
```

#### **Attribute macros**

r[macro.proc.attribute.intro] *Attribute macros* define new <u>outer attributes</u> which can be attached to <u>items</u>, including items in <u>extern blocks</u>, inherent and trait <u>implementations</u>, and <u>trait definitions</u>.

r[macro.proc.attribute.def] Attribute macros are defined by a <u>public</u> <u>function</u> with the proc\_macro\_attribute <u>attribute</u> that has a signature of (TokenStream, TokenStream) -> TokenStream. The first <u>TokenStream</u> is the delimited token tree following the attribute's name, not including the outer delimiters. If the attribute is written as a bare attribute name, the attribute <u>TokenStream</u> is empty. The second <u>TokenStream</u> is the rest of the <u>item</u> including other <u>attributes</u> on the <u>item</u>. The returned <u>TokenStream</u> replaces the <u>item</u> with an arbitrary number of <u>items</u>.

r[macro.proc.attribute.namespace] The proc\_macro\_attribute attribute defines the attribute in the macro namespace in the root of the crate.

For example, this attribute macro takes the input stream and returns it as is, effectively being the no-op of attributes.

```
# #![crate_type = "proc-macro"]
# extern crate proc_macro;
```

```
# use proc_macro::TokenStream;
```

```
#[proc_macro_attribute]
pub fn return_as_is(_attr: TokenStream, item: TokenStream) ->
TokenStream {
    item
}
```

This following example shows the stringified <u>TokenStreams</u> that the attribute macros see. The output will show in the output of the compiler. The output is shown in the comments after the function prefixed with "out:".

```
// my-macro/src/lib.rs
# extern crate proc_macro;
# use proc_macro::TokenStream;
```

```
#[proc_macro_attribute]
pub fn show_streams(attr: TokenStream, item: TokenStream) ->
TokenStream {
    println!("attr: \"{attr}\"");
    println!("item: \"{item}\"");
    item
}
// src/lib.rs
extern crate my_macro;
use my_macro::show_streams;
// Example: Basic function
#[show_streams]
fn invoke1() {}
// out: attr: ""
// out: item: "fn invoke1() {}"
// Example: Attribute with input
#[show_streams(bar)]
fn invoke2() {}
// out: attr: "bar"
// out: item: "fn invoke2() {}"
// Example: Multiple tokens in the input
#[show_streams(multiple => tokens)]
fn invoke3() {}
// out: attr: "multiple => tokens"
// out: item: "fn invoke3() {}"
// Example:
#[show_streams { delimiters }]
fn invoke4() {}
// out: attr: "delimiters"
// out: item: "fn invoke4() {}"
```

r[macro.proc.token]

# Declarative macro tokens and procedural macro tokens

r[macro.proc.token.intro] Declarative macro\_rules macros and procedural macros use similar, but different definitions for tokens (or rather <u>TokenTree s</u>.)

r[macro.proc.token.macro\_rules] Token trees in macro\_rules
(corresponding to tt matchers) are defined as

- Delimited groups ( ( . . . ) , { . . . } , etc)
- All operators supported by the language, both single-character and multi-character ones (+, +=).
  - Note that this set doesn't include the single quote '.
- Literals ("string", 1, etc)
  - Note that negation (e.g. -1) is never a part of such literal tokens, but a separate operator token.
- Identifiers, including keywords (ident, r#ident, fn)
- Lifetimes ('ident)
- Metavariable substitutions in macro\_rules (e.g. \$my\_expr in macro\_rules! mac { (\$my\_expr: expr) => { \$my\_expr } } after the mac's expansion, which will be considered a single token tree regardless of the passed expression)

r[macro.proc.token.tree] Token trees in procedural macros are defined as

- Delimited groups ((...), {...}, etc)
- All punctuation characters used in operators supported by the language (+, but not +=), and also the single quote ' character (typically used in lifetimes, see below for lifetime splitting and joining behavior)
- Literals ("string", 1, etc)

- Negation (e.g. -1) is supported as a part of integer and floating point literals.
- Identifiers, including keywords (ident, r#ident, fn)

r[macro.proc.token.conversion.intro] Mismatches between these two definitions are accounted for when token streams are passed to and from procedural macros.

Note that the conversions below may happen lazily, so they might not happen if the tokens are not actually inspected.

r[macro.proc.token.conversion.to-proc\_macro] When passed to a proc-macro

- All multi-character operators are broken into single characters.
- Lifetimes are broken into a ' character and an identifier.
- The keyword metavariable <u>\$crate</u> is passed as a single identifier.
- All other metavariable substitutions are represented as their underlying token streams.
  - Such token streams may be wrapped into delimited groups (<u>Group</u>) with implicit delimiters (<u>Delimiter::None</u>) when it's necessary for preserving parsing priorities.
  - tt and ident substitutions are never wrapped into such groups and always represented as their underlying token trees.

r[macro.proc.token.conversion.from-proc\_macro] When emitted from a proc macro

- Punctuation characters are glued into multi-character operators when applicable.
- Single quotes ' joined with identifiers are glued into lifetimes.
- Negative literals are converted into two tokens (the and the literal) possibly wrapped into a delimited group (<u>Group</u>) with implicit delimiters (<u>Delimiter::None</u>) when it's necessary for preserving parsing priorities.

r[macro.proc.token.doc-comment] Note that neither declarative nor procedural macros support doc comment tokens (e.g. /// Doc ), so they are

always converted to token streams representing their equivalent #[doc = r"str"] attributes when passed to macros.

r[crate]

# **Crates and source files**

r[crate.syntax] @root Crate -> InnerAttribute\* Item\*

> [!NOTE] Although Rust, like any other language, can be implemented by an interpreter as well as a compiler, the only existing implementation is a compiler, and the language has always been designed to be compiled. For these reasons, this section assumes a compiler.

r[crate.compile-time] Rust's semantics obey a *phase distinction* between compile-time and run-time.<sup>1</sup> Semantic rules that have a *static interpretation* govern the success or failure of compilation, while semantic rules that have a *dynamic interpretation* govern the behavior of the program at run-time.

r[crate.unit] The compilation model centers on artifacts called *crates*. Each compilation processes a single crate in source form, and if successful, produces a single crate in binary form: either an executable or some sort of library.<sup>2</sup>

r[crate.module] A *crate* is a unit of compilation and linking, as well as versioning, distribution, and runtime loading. A crate contains a *tree* of nested <u>module</u> scopes. The top level of this tree is a module that is anonymous (from the point of view of paths within the module) and any item within a crate has a canonical <u>module path</u> denoting its location within the crate's module tree.

r[crate.input-source] The Rust compiler is always invoked with a single source file as input, and always produces a single output crate. The processing of that source file may result in other source files being loaded as modules. Source files have the extension .rs.

r[crate.module-def] A Rust source file describes a module, the name and location of which — in the module tree of the current crate — are defined from outside the source file: either by an explicit [Module][grammar-Module] item in a referencing source file, or by the name of the crate itself.

r[crate.inline-module] Every source file is a module, but not every module needs its own source file: <u>module definitions</u> can be nested within one file.

r[crate.items] Each source file contains a sequence of zero or more [Item] definitions, and may optionally begin with any number of <u>attributes</u> that apply to the containing module, most of which influence the behavior of the compiler.

r[crate.attributes] The anonymous crate module can have additional attributes that apply to the crate as a whole.

[!NOTE] The file's contents may be preceded by a <u>shebang</u>.

```
// Specify the crate name.
#![crate_name = "projx"]
// Specify the type of output artifact.
#![crate_type = "lib"]
// Turn on a warning.
// This can be done in any module, not just the anonymous
crate module.
#![warn(non_camel_case_types)]
```

r[crate.main]

### **Main Functions**

r[crate.main.general] A crate that contains a main <u>function</u> can be compiled to an executable.

r[crate.main.restriction] If a main function is present, it must take no arguments, must not declare any <u>trait or lifetime bounds</u>, must not have any <u>where clauses</u>, and its return type must implement the <u>Termination</u> trait.

```
fn main() {}
fn main() -> ! {
    std::process::exit(0);
}
fn main() -> impl std::process::Termination {
    std::process::ExitCode::SUCCESS
}
```

r[crate.main.import] The main function may be an import, e.g. from an external crate or from the current one.

```
mod foo {
    pub fn bar() {
        println!("Hello, world!");
    }
}
use foo::bar as main;
```

[!NOTE] Types with implementations of <u>Termination</u> in the standard library include:

```
• ()
```

• !

- <u>Infallible</u>
- <u>ExitCode</u>
- Result<T, E> where T: Termination, E: Debug

r[crate.uncaught-foreign-unwinding]

# **Uncaught foreign unwinding**

When a "foreign" unwind (e.g. an exception thrown from C++ code, or a panic! in Rust code using a different panic handler) propagates beyond the main function, the process will be safely terminated. This may take the form of an abort, in which case it is not guaranteed that any Drop calls will be executed, and the error output may be less informative than if the runtime had been terminated by a "native" Rust panic.

For more information, see the <u>panic documentation</u>.

r[crate.no\_main]

### The no\_main attribute

The *no\_main* <u>attribute</u> may be applied at the crate level to disable emitting the main symbol for an executable binary. This is useful when some other object being linked to defines main.

r[crate.crate\_name]

### The crate\_name attribute

r[crate.crate\_name.general] The *crate\_name* <u>attribute</u> may be applied at the crate level to specify the name of the crate with the [MetaNameValueStr] syntax.

```
#![crate_name = "mycrate"]
```

r[crate.crate\_name.restriction] The crate name must not be empty, and must only contain <u>Unicode alphanumeric</u> or  $\_$  (U+005F) characters. 1

This distinction would also exist in an interpreter. Static checks like syntactic analysis, type checking, and lints should happen before the program is executed regardless of when it is executed.

A crate is somewhat analogous to an *assembly* in the ECMA-335 CLI model, a *library* in the SML/NJ Compilation Manager, a *unit* in the Owens and Flatt module system, or a *configuration* in Mesa.

r[cfg]

# **Conditional compilation**

```
r[cfg.syntax]
ConfigurationPredicate ->
      ConfigurationOption
    | ConfigurationAll
    | ConfigurationAny
    | ConfigurationNot
    l `true`
    | `false`
ConfigurationOption ->
    IDENTIFIER ( `=` ( STRING_LITERAL | RAW_STRING_LITERAL ) )?
ConfigurationAll ->
    `all` `(` ConfigurationPredicateList? `)`
ConfigurationAny ->
    `any` `(` ConfigurationPredicateList? `)`
ConfigurationNot ->
    `not` `(` ConfigurationPredicate `)`
ConfigurationPredicateList ->
```

ConfigurationPredicate (`,` ConfigurationPredicate)\* `,`?

r[cfg.general] *Conditionally compiled source code* is source code that is compiled only under certain conditions.

r[cfg.attributes-macro] Source code can be made conditionally compiled using the <u>cfg</u> and <u>cfg\_attr\_attributes</u> and the built-in <u>cfg\_macro</u>.

r[cfg.conditional] Whether to compile can depend on the target architecture of the compiled crate, arbitrary values passed to the compiler, and other things further described below.

r[cfg.predicate] Each form of conditional compilation takes a *configuration predicate* that evaluates to true or false. The predicate is one

of the following:

r[cfg.predicate.option]

A configuration option. The predicate is true if the option is set, and false if it is unset.

r[cfg.predicate.all]

- all() with a comma-separated list of configuration predicates. It is true if all of the given predicates are true, or if the list is empty.
   r[cfg.predicate.any]
- any() with a comma-separated list of configuration predicates. It is true if at least one of the given predicates is true. If there are no predicates, it is false.

r[cfg.predicate.not]

• not() with a configuration predicate. It is true if its predicate is false and false if its predicate is true.

r[cfg.predicate.literal]

• true or false literals, which are always true or false respectively.

r[cfg.option-spec] *Configuration options* are either names or key-value pairs, and are either set or unset.

r[cfg.option-name] Names are written as a single identifier, such as unix.

r[cfg.option-key-value] Key-value pairs are written as an identifier, =, and then a string, such as target\_arch = "x86\_64".

[!NOTE] Whitespace around the = is ignored, so foo="bar" and foo = "bar" are equivalent.

r[cfg.option-key-uniqueness] Keys do not need to be unique. For example, both feature = "std" and feature = "serde" can be set at the same time.

r[cfg.options.set]

# **Set Configuration Options**

r[cfg.options.general] Which configuration options are set is determined statically during the compilation of the crate.

r[cfg.options.target] Some options are *compiler-set* based on data about the compilation.

r[cfg.options.other] Other options are *arbitrarily-set* based on input passed to the compiler outside of the code.

r[cfg.options.crate] It is not possible to set a configuration option from within the source code of the crate being compiled.

[!NOTE] For rustc, arbitrary-set configuration options are set using the <u>--cfg</u> flag. Configuration values for a specified target can be displayed with rustc --print cfg --target \$TARGET.

[!NOTE] Configuration options with the key feature are a convention used by <u>Cargo</u> for specifying compile-time options and optional dependencies.

r[cfg.target\_arch]

#### target\_arch

r[cfg.target\_arch.gen] Key-value option set once with the target's CPU architecture. The value is similar to the first element of the platform's target triple, but not identical.

r[cfg.target\_arch.values] Example values:

```
• "x86"
```

- "x86\_64"
- "mips"
- "powerpc"
- "powerpc64"
- "arm"
- "aarch64"

r[cfg.target\_feature]

#### target\_feature

r[cfg.target\_feature.general] Key-value option set for each platform feature available for the current compilation target.

r[cfg.target\_feature.values] Example values:

- "avx"
- "avx2"
- "crt-static"
- "rdrand"
- "sse"
- "sse2"
- "sse4.1"

See the <u>target feature</u> attribute for more details on the available features.

r[cfg.target\_feature.crt\_static] An additional feature of crt-static is
available to the target\_feature option to indicate that a static C runtime
is available.

r[cfg.target\_os]

#### target\_os

r[cfg.target\_os.general] Key-value option set once with the target's operating system. This value is similar to the second and third element of the platform's target triple.

r[cfg.target\_os.values] Example values:

- "windows"
- "macos"
- "ios"
- "linux"
- "android"
- "freebsd"
- "dragonfly"

- "openbsd"
- "netbsd"
- "none" (typical for embedded targets)

r[cfg.target\_family]

#### target\_family

r[cfg.target\_family.general] Key-value option providing a more generic description of a target, such as the family of the operating systems or architectures that the target generally falls into. Any number of target\_family key-value pairs can be set.

r[cfg.target\_family.values] Example values:

```
• "unix"
```

- "windows"
- "wasm"
- Both "unix" and "wasm"

r[cfg.target\_family.unix]

#### unix and windows

```
unix is set if target_family = "unix" is set.
```

r[cfg.target\_family.windows] windows is set if target\_family =
"windows" is set.

r[cfg.target\_env]

#### target\_env

r[cfg.target\_env.general] Key-value option set with further disambiguating information about the target platform with information about the ABI or libc used. For historical reasons, this value is only defined as not the empty-string when actually needed for disambiguation. Thus, for example, on many GNU platforms, this value will be empty. This value is similar to the fourth element of the platform's target triple. One difference is that embedded ABIs such as gnueabihf will simply define target\_env as "gnu". r[cfg.target\_env.values] Example values:

- ""
- "gnu"
- "msvc"
- "musl"
- "sgx"

r[cfg.target\_abi]

#### target\_abi

r[cfg.target\_abi.general] Key-value option set to further disambiguate the target\_env with information about the target ABI.

r[cfg.target\_abi.disambiguation] For historical reasons, this value is only defined as not the empty-string when actually needed for disambiguation. Thus, for example, on many GNU platforms, this value will be empty.

r[cfg.target\_abi.values] Example values:

- ""
- "llvm"
- "eabihf"
- "abi64"
- "sim"
- "macabi"

r[cfg.target\_endian]

### target\_endian

Key-value option set once with either a value of "little" or "big" depending on the endianness of the target's CPU.

r[cfg.target\_pointer\_width]

#### target\_pointer\_width

r[cfg.target\_pointer\_width.general] Key-value option set once with the target's pointer width in bits.

r[cfg.target\_pointer\_width.values] Example values:

- "16"
- "32"
- "64"

r[cfg.target\_vendor]

#### target\_vendor

r[cfg.target\_vendor.general] Key-value option set once with the vendor of the target.

r[cfg.target\_vendor.values] Example values:

- "apple"
- "fortanix"
- "pc"
- "unknown"

r[cfg.target\_has\_atomic]

#### target\_has\_atomic

r[cfg.target\_has\_atomic.general] Key-value option set for each bit width that the target supports atomic loads, stores, and compare-and-swap operations.

r[cfg.target\_has\_atomic.stdlib] When this cfg is present, all of the stable
[core::sync::atomic] APIs are available for the relevant atomic width.

r[cfg.target\_has\_atomic.values] Possible values:

```
• "8"
```

```
• "16"
```

- "32"
- "64"
- "128"
- "ptr"

r[cfg.test]

#### test

Enabled when compiling the test harness. Done with rustc by using the <u>--test</u> flag. See <u>Testing</u> for more on testing support.

r[cfg.debug\_assertions]

#### debug\_assertions

Enabled by default when compiling without optimizations. This can be used to enable extra debugging code in development but not in production. For example, it controls the behavior of the standard library's [debug\_assert!] macro.

r[cfg.proc\_macro]

#### proc\_macro

Set when the crate being compiled is being compiled with the proc\_macro crate type.

r[cfg.panic]

#### panic

r[cfg.panic.general] Key-value option set depending on the <u>panic</u> <u>strategy</u>. Note that more values may be added in the future.

r[cfg.panic.values] Example values:

- "abort"
- "unwind"

# Forms of conditional compilation

r[cfg.attr]

# The cfg attribute

r[cfg.attr.intro] The *cfg* <u>attribute</u> conditionally includes the form to which it is attached based on a configuration predicate.

```
[!EXAMPLE]
// The function is only included in the build when
compiling for macOS
#[cfg(target_os = "macos")]
fn macos_only() {
 // ...
}
// This function is only included when either foo or bar
is defined
#[cfg(any(foo, bar))]
fn needs_foo_or_bar() {
 // ...
}
// This function is only included when compiling for a
unixish OS with a 32-bit
// architecture
#[cfg(all(unix, target_pointer_width = "32"))]
fn on_32bit_unix() {
 // ...
}
// This function is only included when foo is not defined
#[cfg(not(foo))]
fn needs_not_foo() {
 // ...
```

```
}
// This function is only included when the panic strategy
is set to unwind
#[cfg(panic = "unwind")]
fn when_unwinding() {
   // ...
}
```

r[cfg.attr.syntax] The syntax for the cfg attribute is:

```
@root CfgAttribute -> `cfg` `(` ConfigurationPredicate `)`
```

r[cfg.attr.allowed-positions] The cfg attribute is allowed anywhere attributes are allowed.

r[cfg.attr.duplicates] Multiple cfg attributes may be specified. The form to which the attribute is attached will not be included if any of the cfg predicates are false except as described in [cfg.attr.crate-level-attrs].

r[cfg.attr.effect] If the predicate is true, the form is rewritten to not have the cfg attribute on it. If the predicate is false, the form is removed from the source code.

r[cfg.attr.crate-level-attrs] When a crate-level cfg has a false predicate, the crate itself still exists. Any crate attributes preceding the cfg are kept, and any crate attributes following the cfg are removed as well as removing all of the following crate contents.

```
[!EXAMPLE] The behavior of not removing the preceding
attributes allows you to do things such as include #![no_std] to avoid
linking std even if a #![cfg(...)] has otherwise removed the
contents of the crate. For example:
// This `no_std` attribute is kept even though the crate-
level `cfg`
// attribute is false.
```

```
#![no_std]
```

```
#![cfg(false)]
```

```
// This function is not included.
pub fn example() {}
```

r[cfg.cfg\_attr]

#### The cfg\_attr attribute

r[cfg.cfg\_attr.intro] The *cfg\_attr* <u>attribute</u> conditionally includes attributes based on a configuration predicate.

```
[!EXAMPLE] The following module will either be found at
linux.rs or windows.rs based on the target.
#[ofg_attr(target_os = "linux" _ notb = "linux rs")]
```

```
#[cfg_attr(target_os = "linux", path = "linux.rs")]
#[cfg_attr(windows, path = "windows.rs")]
mod os;
```

r[cfg.cfg\_attr.syntax] The syntax for the cfg\_attr attribute is:

@root CfgAttrAttribute -> `cfg\_attr` `(` ConfigurationPredicate
`,` CfgAttrs? `)`

```
CfgAttrs -> Attr (`,` Attr)* `,`?
```

r[cfg.cfg\_attr.allowed-positions] The cfg\_attr attribute is allowed anywhere attributes are allowed.

r[cfg.cfg\_attr.duplicates] Multiple cfg\_attr attributes may be specified.

r[cfg.cfg\_attr.attr-restriction] The <u>crate type</u> and <u>crate name</u> attributes cannot be used with cfg\_attr.

r[cfg.cfg\_attr.behavior] When the configuration predicate is true, cfg\_attr expands out to the attributes listed after the predicate.

r[cfg.cfg\_attr.attribute-list] Zero, one, or more attributes may be listed. Multiple attributes will each be expanded into separate attributes.

```
[!EXAMPLE]
```

```
#[cfg_attr(feature = "magic", sparkles, crackles)]
fn bewitched() {}
```

// When the `magic` feature flag is enabled, the above will

```
expand to:
#[sparkles]
#[crackles]
fn bewitched() {}
```

```
[!NOTE] The cfg_attr can expand to another cfg_attr. For
example, #[cfg_attr(target_os = "linux", cfg_attr(feature =
"multithreaded", some_other_attribute))] is valid. This example
would be equivalent to #[cfg_attr(all(target_os = "linux",
feature ="multithreaded"), some_other_attribute)].
```

r[cfg.macro]

#### The cfg macro

The built-in cfg macro takes in a single configuration predicate and evaluates to the true literal when the predicate is true and the false literal when it is false.

For example:

```
let machine_kind = if cfg!(unix) {
    "unix"
} else if cfg!(windows) {
    "windows"
} else {
    "unknown"
};
println!("I'm running on a {} machine!", machine_kind);
```

r[items]

# Items

```
r[items.syntax]
Item ->
    OuterAttribute* ( VisItem | MacroItem )
VisItem ->
    Visibility?
    (
        Module
      | ExternCrate
       | UseDeclaration
      | Function
      | TypeAlias
      | Struct
      | Enumeration
      | Union
      | ConstantItem
      | StaticItem
      | Trait
      | Implementation
      | ExternBlock
    )
```

```
MacroItem ->
```

MacroInvocationSemi

| MacroRulesDefinition

r[items.intro] An *item* is a component of a crate. Items are organized within a crate by a nested set of <u>modules</u>. Every crate has a single "outermost" anonymous module; all further items within the crate have <u>paths</u> within the module tree of the crate.

r[items.static-def] Items are entirely determined at compile-time, generally remain fixed during execution, and may reside in read-only memory.

r[items.kinds] There are several kinds of items:

- modules
- extern crate declarations
- <u>use</u> <u>declarations</u>
- <u>function definitions</u>
- <u>type definitions</u>
- <u>struct definitions</u>
- enumeration definitions
- <u>union definitions</u>
- constant items
- <u>static items</u>
- trait definitions
- <u>implementations</u>
- <u>extern blocks</u>

r[items.locations] Items may be declared in the <u>root of the crate</u>, a <u>module</u>, or a <u>block expression</u>.

r[items.associated-locations] A subset of items, called <u>associated items</u>, may be declared in <u>traits</u> and <u>implementations</u>.

r[items.extern-locations] A subset of items, called external items, may be declared in <u>extern blocks</u>.

r[items.decl-order] Items may be defined in any order, with the exception of macro rules which has its own scoping behavior.

r[items.name-resolution] <u>Name resolution</u> of item names allows items to be defined before or after where the item is referred to in the module or block.

See <u>item scopes</u> for information on the scoping rules of items.

r[items.mod]

# Modules

r[items.mod.syntax]

```
Module ->
  `unsafe`? `mod` IDENTIFIER `;`
  | `unsafe`? `mod` IDENTIFIER `{`
    InnerAttribute*
    Item*
    `}`
```

r[items.mod.intro] A module is a container for zero or more <u>items</u>.

r[items.mod.def] A *module item* is a module, surrounded in braces, named, and prefixed with the keyword mod. A module item introduces a new, named module into the tree of modules making up a crate.

r[items.mod.nesting] Modules can nest arbitrarily.

An example of a module:

```
mod math {
    type Complex = (f64, f64);
    fn sin(f: f64) -> f64 {
        /* ... */
        unimplemented!();
#
    }
    fn cos(f: f64) -> f64 {
        /* ... */
        unimplemented!();
#
    }
    fn tan(f: f64) -> f64 {
        /* ... */
        unimplemented!();
#
    }
}
```

r[items.mod.namespace] Modules are defined in the <u>type namespace</u> of the module or block where they are located.

r[items.mod.multiple-items] It is an error to define multiple items with the same name in the same namespace within a module. See the <u>scopes</u> <u>chapter</u> for more details on restrictions and shadowing behavior.

r[items.mod.unsafe] The unsafe keyword is syntactically allowed to appear before the mod keyword, but it is rejected at a semantic level. This allows macros to consume the syntax and make use of the unsafe keyword, before removing it from the token stream.

r[items.mod.outlined]

### **Module Source Filenames**

r[items.mod.outlined.intro] A module without a body is loaded from an external file. When the module does not have a path attribute, the path to the file mirrors the logical <u>module path</u>.

r[items.mod.outlined.search] Ancestor module path components are directories, and the module's contents are in a file with the name of the module plus the .rs extension. For example, the following module structure can have this corresponding filesystem structure:

Module Path	<b>Filesystem Path</b>	<b>File Contents</b>
crate	lib.rs	<pre>mod util;</pre>
crate::util	util.rs	<pre>mod config;</pre>
<pre>crate::util::config</pre>	util/config.rs	

r[items.mod.outlined.search-mod] Module filenames may also be the name of the module as a directory with the contents in a file named mod.rs within that directory. The above example can alternately be expressed with crate::util's contents in a file named util/mod.rs. It is not allowed to have both util.rs and util/mod.rs.

[!NOTE] Prior to rustc 1.30, using mod.rs files was the way to load a module with nested children. It is encouraged to use the new naming convention as it is more consistent, and avoids having many files named mod.rs within a project.

r[items.mod.outlined.path]

#### The path attribute

r[items.mod.outlined.path.intro] The directories and files used for loading external file modules can be influenced with the path attribute.

r[items.mod.outlined.path.search] For path attributes on modules not inside inline module blocks, the file path is relative to the directory the

source file is located. For example, the following code snippet would use the paths shown based on where it is located:

```
#[path = "foo.rs"]
mod c;
```

Source File	c's File Location	c's Module Path
<pre>src/a/b.rs</pre>	<pre>src/a/foo.rs</pre>	<pre>crate::a::b::c</pre>
<pre>src/a/mod.rs</pre>	<pre>src/a/foo.rs</pre>	crate::a::c

r[items.mod.outlined.path.search-nested] For path attributes inside inline module blocks, the relative location of the file path depends on the kind of source file the path attribute is located in. "mod-rs" source files are root modules (such as lib.rs or main.rs) and modules with files named mod.rs. "non-mod-rs" source files are all other module files. Paths for path attributes inside inline module blocks in a mod-rs file are relative to the directory of the mod-rs file including the inline module components as directories. For non-mod-rs files, it is the same except the path starts with a directory with the name of the non-mod-rs module. For example, the following code snippet would use the paths shown based on where it is located:

```
mod inline {
    #[path = "other.rs"]
    mod inner;
}
```

Source File	inner's File Location	inner's Module Path
src/a/b. rs	<pre>src/a/b/inline/oth er.rs</pre>	<pre>crate::a::b::inline:: inner</pre>
src/a/mo d.rs	<pre>src/a/inline/other .rs</pre>	crate::a::inline::in ner

An example of combining the above rules of path attributes on inline modules and nested modules within (applies to both mod-rs and non-mod-rs files):

```
#[path = "thread_files"]
mod thread {
    // Load the `local_data` module from `thread_files/tls.rs`
relative to
    // this source file's directory.
    #[path = "tls.rs"]
    mod local_data;
}
r[items.mod.attributes]
```
#### **Attributes on Modules**

r[items.mod.attributes.intro] Modules, like all items, accept outer attributes. They also accept inner attributes: either after { for a module with a body, or at the beginning of the source file, after the optional BOM and shebang.

r[items.mod.attributes.supported] The built-in attributes that have meaning on a module are cfg, deprecated, doc, the lint check attributes, path, and no implicit prelude. Modules also accept macro attributes.

r[items.extern-crate]

# **Extern crate declarations**

```
r[items.extern-crate.syntax]
ExternCrate -> `extern` `crate` CrateRef AsClause? `;`
CrateRef -> IDENTIFIER | `self`
AsClause -> `as` ( IDENTIFIER | `_` )
```

r[items.extern-crate.intro] An *extern crate declaration* specifies a dependency on an external crate.

r[items.extern-crate.namespace] The external crate is then bound into the declaring scope as the given <u>identifier</u> in the <u>type namespace</u>.

r[items.extern-crate.extern-prelude] Additionally, if the extern crate appears in the crate root, then the crate name is also added to the <u>extern</u> <u>prelude</u>, making it automatically in scope in all modules.

r[items.extern-crate.as] The as clause can be used to bind the imported crate to a different name.

r[items.extern-crate.lookup] The external crate is resolved to a specific soname at compile time, and a runtime linkage requirement to that soname is passed to the linker for loading at runtime. The soname is resolved at compile time by scanning the compiler's library path and matching the optional crate\_name provided against the crate\_name\_attributes that were declared on the external crate when it was compiled. If no crate\_name is provided, a default name\_attribute is assumed, equal to the identifier given in the extern crate\_declaration.

r[items.extern-crate.self] The self crate may be imported which creates a binding to the current crate. In this case the as clause must be used to specify the name to bind it to.

Three examples of extern crate declarations: extern crate pcre;

extern crate std; // equivalent to: extern crate std as std;

extern crate std as ruststd; // linking to 'std' under another
name

r[items.extern-crate.name-restrictions] When naming Rust crates, hyphens are disallowed. However, Cargo packages may make use of them. In such case, when Cargo.toml doesn't specify a crate name, Cargo will transparently replace - with \_ (Refer to <u>RFC 940</u> for more details).

Here is an example:

// Importing the Cargo package hello-world

extern crate hello\_world; // hyphen replaced with an underscore

r[items.extern-crate.underscore]

### **Underscore Imports**

r[items.extern-crate.underscore.intro] An external crate dependency can be declared without binding its name in scope by using an underscore with the form extern crate foo as \_. This may be useful for crates that only need to be linked, but are never referenced, and will avoid being reported as unused.

r[items.extern-crate.underscore.macro\_use] The <u>macro\_use\_attribute</u> works as usual and imports the macro names into the <u>macro\_use\_prelude</u>.

r[items.extern-crate.no\_link]

#### The no\_link attribute

The *no\_link* attribute may be specified on an extern crate item to prevent linking the crate into the output. This is commonly used to load a crate to access only its macros.

r[items.use]

# **Use declarations**

```
r[items.use.syntax]
UseDeclaration -> `use` UseTree `;`
```

```
UseTree ->
    (SimplePath? `::`)? `*`
    | (SimplePath? `::`)? `{` (UseTree ( `,` UseTree )* `,`?)?
`}`
    | SimplePath ( `as` ( IDENTIFIER | `_` ) )?
```

r[items.use.intro] A *use declaration* creates one or more local name bindings synonymous with some other <u>path</u>. Usually a <u>use</u> declaration is used to shorten the path required to refer to a module item. These declarations may appear in <u>modules</u> and <u>blocks</u>, usually at the top. A <u>use</u> declaration is also sometimes called an *import*, or, if it is public, a *re-export*.

r[items.use.forms] Use declarations support a number of convenient shortcuts:

r[items.use.forms.multiple]

 Simultaneously binding a list of paths with a common prefix, using the brace syntax use a::b::{c, d, e::f, g::h::i};

r[items.use.forms.self]

Simultaneously binding a list of paths with a common prefix and their common parent module, using the self keyword, such as use a::b::
 {self, c, d::e};

r[items.use.forms.as]

Rebinding the target name as a new local name, using the syntax use
 p::q::r as x; This can also be used with the last two features: use
 a::b::{self as ab, c as abc}.

r[items.use.forms.glob]

• Binding all paths matching a given prefix, using the asterisk wildcard syntax use a::b::\*;.

r[items.use.forms.nesting]

```
• Nesting groups of the previous features multiple times, such as use
   a::b::{self as ab, c, d::{*, e::f}};
 An example of use declarations:
use std::collections::hash_map::{self, HashMap};
fn foo<T>(_: T){}
        bar(map1:
fn
                        HashMap<String,
                                             usize>,
                                                           map2:
hash_map::HashMap<String, usize>){}
fn main() {
    // use declarations can also exist inside of functions
    use std::option::Option::{Some, None};
                       11
                                                       'foo(vec!
                               Equivalent
                                               to
[std::option::Option::Some(1.0f64),
    // std::option::Option::None]);'
    foo(vec![Some(1.0f64), None]);
    // Both `hash_map` and `HashMap` are in scope.
    let map1 = HashMap::new();
    let map2 = hash_map::HashMap::new();
    bar(map1, map2);
}
```

```
r[items.use.visibility]
```

#### use Visibility

r[items.use.visibility.intro] Like items, use declarations are private to the containing module, by default. Also like items, a use declaration can be public, if qualified by the pub keyword. Such a use declaration serves to *re-export* a name. A public use declaration can therefore *redirect* some public name to a different target definition: even a definition with a private canonical path, inside a different module.

r[items.use.visibility.unambiguous] If a sequence of such redirections form a cycle or cannot be resolved unambiguously, they represent a compile-time error.

An example of re-exporting:

```
mod quux {
    pub use self::foo::{bar, baz};
    pub mod foo {
        pub fn bar() {}
        pub fn baz() {}
    }
}
fn main() {
    quux::bar();
    quux::baz();
}
```

In this example, the module quux re-exports two public names defined in foo.

r[items.use.path]

#### use Paths

r[items.use.path.intro] The <u>paths</u> that are allowed in a <u>use</u> item follow the [SimplePath] grammar and are similar to the paths that may be used in an expression. They may create bindings for:

- Nameable <u>items</u>
- Enum variants
- <u>Built-in types</u>
- <u>Attributes</u>
- <u>Derive macros</u>

r[items.use.path.disallowed] They cannot import <u>associated items</u>, <u>generic parameters</u>, <u>local variables</u>, paths with <u>Self</u>, or <u>tool attributes</u>. More restrictions are described below.

r[items.use.path.namespace] use will create bindings for all <u>namespaces</u> from the imported entities, with the exception that a self import will only import from the type namespace (as described below). For example, the following illustrates creating bindings for the same name in two namespaces:

```
mod stuff {
   pub struct Foo(pub i32);
}
// Imports the `Foo` type and the `Foo` constructor.
use stuff::Foo;
fn example() {
   let ctor = Foo; // Uses `Foo` from the value namespace.
   let x: Foo = ctor(123); // Uses `Foo` From the type
namespace.
}
```

r[items.use.path.edition2018]

[!EDITION-2018] In the 2015 edition, use paths are relative to the crate root. For example:

```
mod foo {
    pub mod example { pub mod iter {} }
    pub mod baz { pub fn foobaz() {} }
}
mod bar {
    // Resolves `foo` from the crate root.
    use foo::example::iter;
    // The `::` prefix explicitly resolves `foo`
    // from the crate root.
    use ::foo::baz::foobaz;
}
```

```
# fn main() {}
```

The 2015 edition does not allow use declarations to reference the <u>extern prelude</u>. Thus, <u>extern crate</u> declarations are still required in 2015 to reference an external crate in a <u>use</u> declaration. Beginning with the 2018 edition, <u>use</u> declarations can specify an external crate dependency the same way <u>extern crate</u> can.

r[items.use.as]

#### as renames

The as keyword can be used to change the name of an imported entity. For example:

```
// Creates a non-public alias `bar` for the function `foo`.
use inner::foo as bar;
mod inner {
   pub fn foo() {}
}
```

r[items.use.multiple-syntax]

#### **Brace** syntax

r[items.use.multiple-syntax.intro] Braces can be used in the last segment of the path to import multiple entities from the previous segment, or, if there are no previous segments, from the current scope. Braces can be nested, creating a tree of paths, where each grouping of segments is logically combined with its parent to create a full path.

```
// Creates bindings to:
// - `std::collections::BTreeSet`
// - `std::collections::hash_map`
// - `std::collections::hash_map::HashMap`
use std::collections::{BTreeSet, hash_map::{self, HashMap}};
```

r[items.use.multiple-syntax.empty] An empty brace does not import anything, though the leading path is validated that it is accessible.

r[items.use.multiple-syntax.edition2018]

[!EDITION-2018] In the 2015 edition, paths are relative to the crate root, so an import such as use {foo, bar}; will import the names foo and bar from the crate root, whereas starting in 2018, those names are relative to the current scope.

r[items.use.self]

#### self imports

r[items.use.self.intro] The keyword self may be used within <u>brace</u> <u>syntax</u> to create a binding of the parent entity under its own name.

```
mod stuff {
    pub fn foo() {}
    pub fn bar() {}
}
mod example {
    // Creates a binding for `stuff` and `foo`.
    use crate::stuff::{self, foo};
    pub fn baz() {
        foo();
        stuff::bar();
    }
}
# fn main() {}
```

r[items.use.self.namespace] self only creates a binding from the <u>type</u> <u>namespace</u> of the parent entity. For example, in the following, only the foo mod is imported:

```
mod bar {
    pub mod foo {}
    pub fn foo() {}
}
// This only imports the module `foo`. The function `foo` lives
in
// the value namespace and is not imported.
use bar::foo::{self};
fn main() {
    foo(); //~ ERROR `foo` is a module
}
```

[!NOTE] self may also be used as the first segment of a path. The usage of self as the first segment and inside a use brace is logically the same; it means the current module of the parent segment, or the current module if there is no parent segment. See <u>self</u> in the paths chapter for more information on the meaning of a leading self.

r[items.use.glob]

### **Glob** imports

r[items.use.glob.intro] The \* character may be used as the last segment of a use path to import all importable entities from the entity of the preceding segment. For example:

```
// Creates a non-public alias to `bar`.
use foo::*;
mod foo {
    fn i_am_private() {}
    enum Example {
        V1,
        V2,
    }
    pub fn bar() {
        // Creates local aliases to `V1` and `V2`
        // of the `Example` enum.
        use Example::*;
        let x = V1;
    }
}
```

r[items.use.glob.shadowing] Items and named imports are allowed to shadow names from glob imports in the same <u>namespace</u>. That is, if there is a name already defined by another item in the same namespace, the glob import will be shadowed. For example:

```
// This creates a binding to the `clashing::Foo` tuple struct
// constructor, but does not import its type because that
would
// conflict with the `Foo` struct defined here.
//
// Note that the order of definition here is unimportant.
use clashing::*;
struct Foo {
    field: f32,
```

```
fn do_stuff() {
    // Uses the constructor from `clashing::Foo`.
    let f1 = Foo(123);
    // The struct expression uses the type from
    // the `Foo` struct defined above.
    let f2 = Foo { field: 1.0 };
    // `Bar` is also in scope due to the glob import.
    let z = Bar {};
}
mod clashing {
    pub struct Foo(pub i32);
    pub struct Bar {}
}
```

r[items.use.glob.last-segment-only] \* cannot be used as the first or intermediate segments.

r[items.use.glob.self-import] \* cannot be used to import a module's
contents into itself (such as use self::\*;).

r[items.use.glob.edition2018]

}

[!EDITION-2018] In the 2015 edition, paths are relative to the crate root, so an import such as use \*; is valid, and it means to import everything from the crate root. This cannot be used in the crate root itself.

r[items.use.as-underscore]

# **Underscore Imports**

r[items.use.as-underscore.intro] Items can be imported without binding to a name by using an underscore with the form use path as \_\_. This is particularly useful to import a trait so that its methods may be used without importing the trait's symbol, for example if the trait's symbol may conflict with another symbol. Another example is to link an external crate without importing its name.

r[items.use.as-underscore.glob] Asterisk glob imports will import items imported with \_\_\_\_ in their unnameable form.

```
mod foo {
    pub trait Zoo {
        fn zoo(&self) {}
    }
    impl<T> Zoo for T {}
}
use self::foo::Zoo as _;
struct Zoo; // Underscore import avoids name conflict with
this item.
fn main() {
    let z = Zoo;
    z.zoo();
}
```

r[items.use.as-underscore.macro] The unique, unnameable symbols are created after macro expansion so that macros may safely emit multiple references to \_\_\_\_\_ imports. For example, the following should not produce an error:

```
macro_rules! m {
   ($item: item) => { $item $item }
}
```

```
m!(use std as _;);
// This expands to:
// use std as _;
// use std as _;
```

r[items.use.restrictions]

# Restrictions

The following are restrictions for valid use declarations: r[items.use.restrictions.crate]

• use crate; must use as to define the name to which to bind the crate root.

r[items.use.restrictions.self]

use {self}; is an error; there must be a leading segment when using self.

r[items.use.restrictions.duplicate-name]

• As with any item definition, use imports cannot create duplicate bindings of the same name in the same namespace in a module or block.

r[items.use.restrictions.macro-crate]

- use paths with \$crate are not allowed in a <u>macro rules</u> expansion. r[items.use.restrictions.variant]
- use paths cannot refer to enum variants through a <u>type alias</u>. For example:

   enum MyEnum {
   MyVariant
   type TypeAlias = MyEnum;

   use MyEnum::MyVariant; //~ OK
   use TypeAlias::MyVariant; //~ ERROR

   r[items.use.ambiguities]

# Ambiguities

[!NOTE] This section is incomplete.

r[items.use.ambiguities.intro] Some situations are an error when there is an ambiguity as to which name a use declaration refers. This happens when there are two name candidates that do not resolve to the same entity.

r[items.use.ambiguities.glob] Glob imports are allowed to import conflicting names in the same namespace as long as the name is not used. For example:

```
mod foo {
    pub struct Qux;
}
mod bar {
    pub struct Qux;
}
use foo::*;
use bar::*; //~ OK, no name conflict.
fn main() {
    // This would be an error, due to the ambiguity.
    //let x = Qux;
}
```

Multiple glob imports are allowed to import the same name, and that name is allowed to be used, if the imports are of the same item (following re-exports). The visibility of the name is the maximum visibility of the imports. For example:

```
mod foo {
    pub struct Qux;
}
mod bar {
```

```
pub use super::foo::Qux;
}
// These both import the same `Qux`. The visibility of `Qux`
// is `pub` because that is the maximum visibility between
// these two `use` declarations.
pub use bar::*;
use foo::*;
fn main() {
    let _: Qux = Qux;
}
```

r[items.fn]

# **Functions**

```
r[items.fn.syntax]
Function ->
    FunctionQualifiers `fn` IDENTIFIER GenericParams?
        `(` FunctionParameters? `)`
        FunctionReturnType? WhereClause?
        ( BlockExpression | `;` )
FunctionQualifiers -> `const`? `async`?[^async-edition]
ItemSafety?[^extern-gualifiers] (`extern` Abi?)?
ItemSafety -> `safe`[^extern-safe] | `unsafe`
Abi -> STRING LITERAL | RAW STRING LITERAL
FunctionParameters ->
      SelfParam `,`?
    | (SelfParam `,`)? FunctionParam (`,` FunctionParam)* `,`?
SelfParam -> OuterAttribute* ( ShorthandSelf | TypedSelf )
ShorthandSelf -> (`&` | `&` Lifetime)? `mut`? `self`
TypedSelf -> `mut`? `self` `:` Type
FunctionParam -> OuterAttribute* ( FunctionParamPattern | `...`
Type[^fn-param-2015] )
FunctionParamPattern -> PatternNoTopAlt `:` ( Type | `...` )
FunctionReturnType -> `->` Type
1
```

The async qualifier is not allowed in the 2015 edition.

The safe function qualifier is only allowed semantically within extern blocks.

3

Relevant to editions earlier than Rust 2024: Within extern blocks, the safe or unsafe function qualifier is only allowed when the extern is qualified as unsafe.

4

Function parameters with only a type are only allowed in an associated function of a <u>trait item</u> in the 2015 edition.

r[items.fn.intro] A *function* consists of a <u>block</u> (that's the *body* of the function), along with a name, a set of parameters, and an output type. Other than a name, all these are optional.

r[items.fn.namespace] Functions are declared with the keyword fn which defines the given name in the <u>value namespace</u> of the module or block where it is located.

r[items.fn.signature] Functions may declare a set of *input <u>variables</u>* as parameters, through which the caller passes arguments into the function, and the *output <u>type</u>* of the value the function will return to its caller on completion.

r[items.fn.implicit-return] If the output type is not explicitly stated, it is the <u>unit type</u>.

r[items.fn.fn-item-type] When referred to, a *function* yields a first-class *value* of the corresponding zero-sized *function item type*, which when called evaluates to a direct call to the function.

For example, this is a simple function:

```
fn answer_to_life_the_universe_and_everything() -> i32 {
    return 42;
```

```
}
```

r[items.fn.safety-qualifiers] The safe function is semantically only allowed when used in an <u>extern block</u>.

```
r[items.fn.params]
```

#### **Function parameters**

r[items.fn.params.intro] Function parameters are irrefutable <u>patterns</u>, so any pattern that is valid in an else-less <u>let</u> binding is also valid as a parameter:

```
fn first((value, _): (i32, i32)) -> i32 { value }
```

r[items.fn.params.self-pat] If the first parameter is a [SelfParam], this indicates that the function is a <u>method</u>.

r[items.fn.params.self-restriction] Functions with a self parameter may only appear as an <u>associated function</u> in a <u>trait</u> or <u>implementation</u>.

r[items.fn.params.varargs] A parameter with the ... token indicates a <u>variadic function</u>, and may only be used as the last parameter of an <u>external</u> <u>block</u> function. The variadic parameter may have an optional identifier, such as args: ....

r[items.fn.body]

# **Function body**

r[items.fn.body.intro] The body block of a function is conceptually wrapped in another block that first binds the argument patterns and then **returns** the value of the function's body. This means that the tail expression of the block, if evaluated, ends up being returned to the caller. As usual, an explicit return expression within the body of the function will short-cut that implicit return, if reached.

For example, the function above behaves as if it was written as:

```
// argument_0 is the actual first argument passed from the
caller
```

```
let (value, _) = argument_0;
return {
    value
}
```

};

r[items.fn.body.bodyless] Functions without a body block are terminated with a semicolon. This form may only appear in a <u>trait</u> or <u>external block</u>.

r[items.fn.generics]

### **Generic functions**

r[items.fn.generics.intro] A *generic function* allows one or more *parameterized types* to appear in its signature. Each type parameter must be explicitly declared in an angle-bracket-enclosed and comma-separated list, following the function name.

```
// foo is generic over A and B
fn foo<A, B>(x: A, y: B) {
# }
```

r[items.fn.generics.param-names] Inside the function signature and body, the name of the type parameter can be used as a type name.

r[items.fn.generics.param-bounds] <u>Trait</u> bounds can be specified for type parameters to allow methods with that trait to be called on values of that type. This is specified using the where syntax:

```
# use std::fmt::Debug;
fn foo<T>(x: T) where T: Debug {
# }
```

r[items.fn.generics.mono] When a generic function is referenced, its type is instantiated based on the context of the reference. For example, calling the foo function here:

```
use std::fmt::Debug;
fn foo<T>(x: &[T]) where T: Debug {
    // details elided
}
foo(&[1, 2]);
```

will instantiate type parameter T with i32.

r[items.fn.generics.explicit-arguments] The type parameters can also be explicitly supplied in a trailing <u>path</u> component after the function name. This might be necessary if there is not sufficient context to determine the type parameters. For example,  $mem::size_of::<u32>() == 4$ .

r[items.fn.extern]

# **Extern function qualifier**

r[items.fn.extern.intro] The extern function qualifier allows providing function *definitions* that can be called with a particular ABI:

extern "ABI" fn foo() { /\* ... \*/ }

r[items.fn.extern.def] These are often used in combination with <u>external</u> <u>block</u> items which provide function *declarations* that can be used to call functions without providing their *definition*:

```
unsafe extern "ABI" {
   unsafe fn foo(); /* no body */
   safe fn bar(); /* no body */
}
unsafe { foo() };
bar();
```

r[items.fn.extern.default-abi] When "extern" Abi?\* is omitted from
FunctionQualifiers in function items, the ABI "Rust" is assigned. For
example:

```
fn foo() {}
```

is equivalent to:

```
extern "Rust" fn foo() {}
```

r[items.fn.extern.foreign-call] Functions can be called by foreign code, and using an ABI that differs from Rust allows, for example, to provide functions that can be called from other programming languages like C:

```
// Declares a function with the "C" ABI
extern "C" fn new_i32() -> i32 { 0 }
// Declares a function with the "stdcall" ABI
# #[cfg(any(windows, target_arch = "x86"))]
extern "stdcall" fn new_i32_stdcall() -> i32 { 0 }
```

r[items.fn.extern.default-extern] Just as with <u>external block</u>, when the extern keyword is used and the "ABI" is omitted, the ABI used defaults to "C". That is, this:

extern fn new\_i32() -> i32 { 0 }
let fptr: extern fn() -> i32 = new\_i32;
is equivalent to:
extern "C" fn new\_i32() -> i32 { 0 }
let fptr: extern "C" fn() -> i32 = new\_i32;

r[items.fn.extern.unwind]

# Unwinding

r[items.fn.extern.unwind.intro] Most ABI strings come in two variants, one with an -unwind suffix and one without. The Rust ABI always permits unwinding, so there is no Rust-unwind ABI. The choice of ABI, together with the runtime <u>panic handler</u>, determines the behavior when unwinding out of a function.

r[items.fn.extern.unwind.behavior] The table below indicates the behavior of an unwinding operation reaching each type of ABI boundary (function declaration or definition using the corresponding ABI string). Note that the Rust runtime is not affected by, and cannot have an effect on, any unwinding that occurs entirely within another language's runtime, that is, unwinds that are thrown and caught without reaching a Rust ABI boundary.

The panic -unwind column refers to <u>panicking</u> via the <u>panic</u>! macro and similar standard library mechanisms, as well as to any other Rust operations that cause a panic, such as out-of-bounds array indexing or integer overflow.

The "unwinding" ABI category refers to "Rust" (the implicit ABI of Rust functions not marked extern), "C-unwind", and any other ABI with -unwind in its name. The "non-unwinding" ABI category refers to all other ABI strings, including "C" and "stdcall".

Native unwinding is defined per-target. On targets that support throwing and catching C++ exceptions, it refers to the mechanism used to implement this feature. Some platforms implement a form of unwinding referred to as <u>"forced unwinding"</u>; longjmp on Windows and pthread\_exit in glibc

are implemented this way. Forced unwinding is explicitly excluded from the "Native unwind" column in the table.

panic runtime	ABI	panic-unwind	Native unwind (unforced)
panic= unwind	unwinding	unwind	unwind
panic= unwind	non- unwinding	abort (see notes below)	<u>undefined</u> <u>behavior</u>
panic= abort	unwinding	panic aborts without unwinding	abort
panic= abort	non- unwinding	panic aborts without unwinding	<u>undefined</u> <u>behavior</u>

r[items.fn.extern.abort] With panic=unwind, when a panic is turned into an abort by a non-unwinding ABI boundary, either no destructors (Drop calls) will run, or all destructors up until the ABI boundary will run. It is unspecified which of those two behaviors will happen.

For other considerations and limitations regarding unwinding across FFI boundaries, see the <u>relevant section in the Panic documentation</u>.

r[items.fn.const]

# **Const functions**

r[items.fn.const.intro] Functions qualified with the const keyword are <u>const functions</u>, as are <u>tuple struct</u> and <u>tuple variant</u> constructors. *Const functions* can be called from within <u>const contexts</u>.

r[items.fn.const.extern] Const functions may use the  $\underline{\mathsf{extern}}$  function qualifier.

r[items.fn.const.exclusivity] Const functions are not allowed to be <u>async</u>. r[items.fn.async]

# **Async functions**

r[items.fn.async.intro] Functions may be qualified as async, and this can also be combined with the unsafe qualifier:

```
async fn regular_example() { }
async unsafe fn unsafe_example() { }
```

r[items.fn.async.future] Async functions do no work when called: instead, they capture their arguments into a future. When polled, that future will execute the function's body.

r[items.fn.async.desugar-brief] An async function is roughly equivalent to a function that returns <u>impl Future</u> and with an <u>async move block</u> as its body:

```
// Source
async fn example(x: &str) -> usize {
    x.len()
}
is roughly equivalent to:
# use std::future::Future;
// Desugared
fn example<'a>(x: &'a str) -> impl Future<Output = usize> + 'a
{
    async move { x.len() }
}
```

r[items.fn.async.desugar] The actual desugaring is more complex: r[items.fn.async.lifetime-capture]

• The return type in the desugaring is assumed to capture all lifetime parameters from the async fn declaration. This can be seen in the desugared example above, which explicitly outlives, and hence captures, 'a.

r[items.fn.async.param-capture]

• The <u>async move block</u> in the body captures all function parameters, including those that are unused or bound to a \_ pattern. This ensures that function parameters are dropped in the same order as they would be if the function were not async, except that the drop occurs when the returned future has been fully awaited.

For more information on the effect of async, see <u>async blocks</u>.

```
r[items.fn.async.edition2018]
```

[!EDITION-2018] Async functions are only available beginning with Rust 2018.

r[items.fn.async.safety]

# Combining async and unsafe

r[items.fn.async.safety.intro] It is legal to declare a function that is both async and unsafe. The resulting function is unsafe to call and (like any async function) returns a future. This future is just an ordinary future and thus an unsafe context is not required to "await" it:

```
// Returns a future that, when awaited, dereferences x.
11
// Soundness condition: `x` must be safe to dereference until
// the resulting future is complete.
async unsafe fn unsafe_example(x: *const i32) -> i32 {
  *х
}
async fn safe_example() {
     // An `unsafe` block is required to invoke the function
initially:
    let p = 22;
    let future = unsafe { unsafe_example(&p) };
    // But no `unsafe` block required here. This will
    // read the value of `p`:
    let q = future.await;
}
```
Note that this behavior is a consequence of the desugaring to a function that returns an *impl Future* -- in this case, the function we desugar to is an unsafe function, but the return value remains the same.

Unsafe is used on an async function in precisely the same way that it is used on other functions: it indicates that the function imposes some additional obligations on its caller to ensure soundness. As in any other unsafe function, these conditions may extend beyond the initial call itself -- in the snippet above, for example, the unsafe\_example function took a pointer  $\times$  as argument, and then (when awaited) dereferenced that pointer. This implies that  $\times$  would have to be valid until the future is finished executing, and it is the caller's responsibility to ensure that.

r[items.fn.attributes]

#### **Attributes on functions**

r[items.fn.attributes.intro] <u>Outer attributes</u> are allowed on functions. <u>Inner attributes</u> are allowed directly after the { inside its body <u>block</u>.

This example shows an inner attribute on a function. The function is documented with just the word "Example".

```
fn documented() {
    #![doc = "Example"]
}
```

[!NOTE] Except for lints, it is idiomatic to only use outer attributes on function items.

r[items.fn.attributes.builtin-attributes] The attributes that have meaning on a function are <u>cfg</u>, <u>cfg attr</u>, <u>deprecated</u>, <u>doc</u>, <u>export name</u>, <u>link section</u>, <u>no mangle</u>, <u>the lint check attributes</u>, <u>must use</u>, <u>the</u> <u>procedural macro attributes</u>, <u>the testing attributes</u>, and <u>the optimization hint</u> <u>attributes</u>. Functions also accept attributes macros.

```
r[items.fn.param-attributes]
```

#### **Attributes on function parameters**

r[items.fn.param-attributes.intro] <u>Outer attributes</u> are allowed on function parameters and the permitted <u>built-in attributes</u> are restricted to cfg, cfg\_attr, allow, warn, deny, and forbid.

```
fn len(
    #[cfg(windows)] slice: &[u16],
    #[cfg(not(windows))] slice: &[u8],
) -> usize {
    slice.len()
}
```

r[items.fn.param-attributes.parsed-attributes] Inert helper attributes used by procedural macro attributes applied to items are also allowed but be careful to not include these inert attributes in your final TokenStream.

```
For example, the following code defines an inert
some_inert_attribute attribute that is not formally defined anywhere and
the some_proc_macro_attribute procedural macro is responsible for
detecting its presence and removing it from the output token stream.
#[some_proc_macro_attribute]
fn foo_oof(#[some_inert_attribute] arg: u8) {
}
```

r[items.type]

# **Type aliases**

r[items.type.syntax]

TypeAlias ->

```
`type` IDENTIFIER GenericParams? ( `:` TypeParamBounds )?
WhereClause?
```

```
( `=` Type WhereClause?)? `;`
```

r[items.type.intro] A *type alias* defines a new name for an existing <u>type</u> in the <u>type namespace</u> of the module or block where it is located. Type aliases are declared with the keyword type. Every value has a single, specific type, but may implement several different traits, and may be compatible with several different type constraints.

For example, the following defines the type Point as a synonym for the type (u8, u8), the type of pairs of unsigned 8 bit integers:

type Point = (u8, u8); let p: Point = (41, 68);

r[items.type.constructor-alias] A type alias to a tuple-struct or unit-struct cannot be used to qualify that type's constructor:

```
struct MyStruct(u32);
```

```
use MyStruct as UseAlias;
type TypeAlias = MyStruct;
```

```
let _ = UseAlias(5); // OK
let _ = TypeAlias(5); // Doesn't work
```

r[items.type.associated-type] A type alias, when not used as an <u>associated type</u>, must include a [Type][grammar-Type] and may not include [TypeParamBounds].

r[items.type.associated-trait] A type alias, when used as an <u>associated</u> <u>type</u> in a <u>trait</u>, must not include a [Type][grammar-Type] specification but may include [TypeParamBounds].

r[items.type.associated-impl] A type alias, when used as an <u>associated</u> <u>type</u> in a <u>trait impl</u>, must include a [Type][grammar-Type] specification and may not include [TypeParamBounds].

r[items.type.deprecated] Where clauses before the equals sign on a type alias in a trait impl (like type TypeAlias<T> where T: Foo = Bar<T>) are deprecated. Where clauses after the equals sign (like type TypeAlias<T> = Bar<T> where T: Foo) are preferred. r[items.struct]

## Structs

```
r[items.struct.syntax]
```

```
Struct ->
```

StructStruct

| TupleStruct

```
StructStruct ->
```

```
`struct` IDENTIFIER GenericParams? WhereClause? ( `{`
StructFields? `}` | `;` )
```

TupleStruct ->

```
`struct` IDENTIFIER GenericParams? `(` TupleFields? `)`
WhereClause? `;`
```

```
StructFields -> StructField (`,` StructField)* `,`?
```

StructField -> OuterAttribute\* Visibility? IDENTIFIER `:` Type

```
TupleFields -> TupleField (`,` TupleField)* `,`?
```

TupleField -> OuterAttribute\* Visibility? Type

r[items.struct.intro] A *struct* is a nominal <u>struct type</u> defined with the keyword <u>struct</u>.

r[items.struct.namespace] A struct declaration defines the given name in the <u>type namespace</u> of the module or block where it is located.

An example of a struct item and its use:

```
struct Point {x: i32, y: i32}
let p = Point {x: 10, y: 11};
let px: i32 = p.x;
```

r[items.struct.tuple] A *tuple struct* is a nominal <u>tuple type</u>, and is also defined with the keyword struct. In addition to defining a type, it also defines a constructor of the same name in the <u>value namespace</u>. The

constructor is a function which can be called to create a new instance of the struct. For example:

```
struct Point(i32, i32);
let p = Point(10, 11);
let px: i32 = match p { Point(x, _) => x };
```

r[items.struct.unit] A *unit-like struct* is a struct without any fields, defined by leaving off the list of fields entirely. Such a struct implicitly defines a <u>constant</u> of its type with the same name. For example:

```
struct Cookie;
let c = [Cookie, Cookie {}, Cookie, Cookie {}];
is equivalent to
```

```
struct Cookie {}
const Cookie: Cookie = Cookie {};
let c = [Cookie, Cookie {}, Cookie, Cookie {}];
```

r[items.struct.layout] The precise memory layout of a struct is not specified. One can specify a particular layout using the <u>repr\_attribute</u>.

r[items.enum]

## Enumerations

r[items.enum.syntax]

```
Enumeration ->
```

`enum` IDENTIFIER GenericParams? WhereClause? `{`
EnumVariants? `}`

```
EnumVariants -> EnumVariant ( `,` EnumVariant )* `,`?
```

EnumVariant ->

```
OuterAttribute* Visibility?
```

```
IDENTIFIER ( EnumVariantTuple | EnumVariantStruct )?
EnumVariantDiscriminant?
```

```
EnumVariantTuple -> `(` TupleFields? `)`
```

```
EnumVariantStruct -> `{` StructFields? `}`
```

```
EnumVariantDiscriminant -> `=` Expression
```

r[items.enum.intro] An *enumeration*, also referred to as an *enum*, is a simultaneous definition of a nominal <u>enumerated type</u> as well as a set of *constructors*, that can be used to create or pattern-match values of the corresponding enumerated type.

r[items.enum.decl] Enumerations are declared with the keyword enum.

r[items.enum.namespace] The enum declaration defines the enumeration type in the <u>type namespace</u> of the module or block where it is located.

An example of an enum item and its use:

```
enum Animal {
Dog,
Cat,
}
```

```
let mut a: Animal = Animal::Dog;
a = Animal::Cat;
```

r[items.enum.constructor] Enum constructors can have either named or unnamed fields:

```
enum Animal {
    Dog(String, f64),
    Cat { name: String, weight: f64 },
}
let mut a: Animal = Animal::Dog("Cocoa".to_string(), 37.2);
a = Animal::Cat { name: "Spotty".to_string(), weight: 2.7 };
```

In this example, Cat is a *struct-like enum variant*, whereas Dog is simply called an enum variant.

r[items.enum.fieldless] An enum where no constructors contain fields are called a *field-less enum*. For example, this is a fieldless enum:

```
enum Fieldless {
    Tuple(),
    Struct{},
    Unit,
}
```

r[items.enum.unit-only] If a field-less enum only contains unit variants, the enum is called an *unit-only enum*. For example:

```
enum Enum {
Foo = 3,
Bar = 2,
Baz = 1,
}
```

r[items.enum.constructor-names] Variant constructors are similar to <u>struct</u> definitions, and can be referenced by a path from the enumeration name, including in <u>use declarations</u>.

r[items.enum.constructor-namespace] Each variant defines its type in the <u>type namespace</u>, though that type cannot be used as a type specifier. Tuplelike and unit-like variants also define a constructor in the <u>value namespace</u>. r[items.enum.struct-expr] A struct-like variant can be instantiated with a <u>struct expression</u>.

r[items.enum.tuple-expr] A tuple-like variant can be instantiated with a <u>call expression</u> or a <u>struct expression</u>.

r[items.enum.path-expr] A unit-like variant can be instantiated with a <u>path expression</u> or a <u>struct expression</u>. For example:

```
enum Examples {
    UnitLike,
    TupleLike(i32),
    StructLike { value: i32 },
}
use Examples::*; // Creates aliases to all variants.
let x = UnitLike; // Path expression of the const item.
let x = UnitLike {}; // Struct expression.
let y = TupleLike(123); // Call expression.
let y = TupleLike { 0: 123 }; // Struct expression using
integer field names.
let z = StructLike { value: 123 }; // Struct expression.
```

r[items.enum.discriminant]

## Discriminants

r[items.enum.discriminant.intro] Each enum instance has a *discriminant*: an integer logically associated to it that is used to determine which variant it holds.

r[items.enum.discriminant.repr-rust] Under the <u>Rust representation</u>, the discriminant is interpreted as an <u>isize</u> value. However, the compiler is allowed to use a smaller type (or another means of distinguishing variants) in its actual memory layout.

## Assigning discriminant values

r[items.enum.discriminant.explicit]

#### **Explicit discriminants**

r[items.enum.discriminant.explicit.intro] In two circumstances, the discriminant of a variant may be explicitly set by following the variant name with = and a <u>constant expression</u>:

r[items.enum.discriminant.explicit.unit-only]

1. if the enumeration is "<u>unit-only</u>".

r[items.enum.discriminant.explicit.primitive-repr] 2. if a <u>primitive</u> <u>representation</u> is used. For example:

```
#[repr(u8)]
enum Enum {
    Unit = 3,
    Tuple(u16),
    Struct {
        a: u8,
        b: u16,
    } = 1,
}
```

r[items.enum.discriminant.implicit]

#### **Implicit discriminants**

If a discriminant for a variant is not specified, then it is set to one higher than the discriminant of the previous variant in the declaration. If the discriminant of the first variant in the declaration is unspecified, then it is set to zero.

r[items.enum.discriminant.restrictions]

#### Restrictions

r[items.enum.discriminant.restrictions.same-discriminant] It is an error when two variants share the same discriminant.

```
enum SharedDiscriminantError {
   SharedA = 1,
   SharedB = 1
}
enum SharedDiscriminantError2 {
   Zero, // 0
   One, // 1
   OneToo = 1 // 1 (collision with previous!)
}
```

r[items.enum.discriminant.restrictions.above-max-discriminant] It is also an error to have an unspecified discriminant where the previous discriminant is the maximum value for the size of the discriminant.

```
#[repr(u8)]
enum OverflowingDiscriminantError {
    Max = 255,
    MaxPlusOne // Would be 256, but that overflows the enum.
}
```

```
#[repr(u8)]
enum OverflowingDiscriminantError2 {
    MaxMinusOne = 254, // 254
    Max, // 255
    MaxPlusOne // Would be 256, but that overflows the
enum.
}
```

### **Accessing discriminant**

#### Via mem::discriminant

r[items.enum.discriminant.access-opaque]

[std::mem::discriminant] returns an opaque reference to the discriminant of an enum value which can be compared. This cannot be used to get the value of the discriminant.

r[items.enum.discriminant.coercion]

#### Casting

r[items.enum.discriminant.coercion.intro] If an enumeration is <u>unit-only</u> (with no tuple and struct variants), then its discriminant can be directly accessed with a <u>numeric cast</u>; e.g.:

```
enum Enum {
    Foo,
    Bar,
    Baz,
}
assert_eq!(0, Enum::Foo as isize);
assert_eq!(1, Enum::Bar as isize);
assert_eq!(2, Enum::Baz as isize);
```

r[items.enum.discriminant.coercion.fieldless] <u>Field-less enums</u> can be casted if they do not have explicit discriminants, or where only unit variants are explicit.

```
enum Fieldless {
    Tuple(),
    Struct{},
    Unit,
}
assert_eq!(0, Fieldless::Tuple() as isize);
assert_eq!(1, Fieldless::Struct{} as isize);
assert_eq!(2, Fieldless::Unit as isize);
#[repr(u8)]
enum FieldlessWithDiscriminants {
    First = 10,
    Tuple(),
    Second = 20,
    Struct{},
    Unit,
}
assert_eq!(10, FieldlessWithDiscriminants::First as u8);
assert_eq!(11, FieldlessWithDiscriminants::Tuple() as u8);
assert_eq!(20, FieldlessWithDiscriminants::Second as u8);
assert_eq!(21, FieldlessWithDiscriminants::Struct{} as u8);
assert_eq!(22, FieldlessWithDiscriminants::Unit as u8);
```

#### **Pointer casting**

r[items.enum.discriminant.access-memory]

If the enumeration specifies a <u>primitive representation</u>, then the discriminant may be reliably accessed via unsafe pointer casting:

```
#[repr(u8)]
enum Enum {
    Unit,
    Tuple(bool),
    Struct{a: bool},
}
```

```
impl Enum {
    fn discriminant(&self) -> u8 {
        unsafe { *(self as *const Self as *const u8) }
    }
}
let unit_like = Enum::Unit;
let tuple_like = Enum::Tuple(true);
let struct_like = Enum::Struct{a: false};
assert_eq!(0, unit_like.discriminant());
assert_eq!(1, tuple_like.discriminant());
r[items.enum.empty]
```

#### **Zero-variant enums**

r[items.enum.empty.intro] Enums with zero variants are known as *zero-variant enums*. As they have no valid values, they cannot be instantiated.

```
enum ZeroVariants {}
```

r[items.enum.empty.uninhabited] Zero-variant enums are equivalent to the <u>never type</u>, but they cannot be coerced into other types.

```
# enum ZeroVariants {}
let x: ZeroVariants = panic!();
let y: u32 = x; // mismatched type error
```

r[items.enum.variant-visibility]

### **Variant visibility**

Enum variants syntactically allow a [Visibility] annotation, but this is rejected when the enum is validated. This allows items to be parsed with a unified syntax across different contexts where they are used.

```
macro_rules! mac_variant {
    ($vis:vis $name:ident) => {
        enum $name {
            $vis Unit,
            $vis Tuple(u8, u16),
            $vis Struct { f: u8 },
        }
    }
}
// Empty `vis` is allowed.
mac_variant! { E }
11
    This
              allowed,
          is
                         since it is removed
                                                  before
                                                          being
validated.
#[cfg(false)]
enum E {
    pub U,
    pub(crate) T(u8),
    pub(super) T { f: String }
}
```

r[items.union]

## Unions

r[items.union.syntax]

Union ->

`union` IDENTIFIER GenericParams? WhereClause? `{`
StructFields? `}`

r[items.union.intro] A union declaration uses the same syntax as a struct declaration, except with union in place of struct.

r[items.union.namespace] A union declaration defines the given name in the <u>type namespace</u> of the module or block where it is located.

```
#[repr(C)]
union MyUnion {
    f1: u32,
    f2: f32,
}
```

r[items.union.common-storage] The key property of unions is that all fields of a union share common storage. As a result, writes to one field of a union can overwrite its other fields, and size of a union is determined by the size of its largest field.

r[items.union.field-restrictions] Union field types are restricted to the following subset of types:

r[items.union.field-copy]

• Copy types

r[items.union.field-references]

• References (&T and &mut T for arbitrary T)

r[items.union.field-manually-drop]

• ManuallyDrop<T> (for arbitrary T)

r[items.union.field-tuple]

• Tuples and arrays containing only allowed union field types

r[items.union.drop] This restriction ensures, in particular, that union fields never need to be dropped. Like for structs and enums, it is possible to impl Drop for a union to manually define what happens when it gets dropped.

r[items.union.fieldless] Unions without any fields are not accepted by the compiler, but can be accepted by macros.

r[items.union.init]

### Initialization of a union

r[items.union.init.intro] A value of a union type can be created using the same syntax that is used for struct types, except that it must specify exactly one field:

```
# union MyUnion { f1: u32, f2: f32 }
#
let u = MyUnion { f1: 1 };
```

r[items.union.init.result] The expression above creates a value of type MyUnion and initializes the storage using field f1. The union can be accessed using the same syntax as struct fields:

```
# union MyUnion { f1: u32, f2: f32 }
#
# let u = MyUnion { f1: 1 };
let f = unsafe { u.f1 };
```

r[items.union.fields]

## **Reading and writing union fields**

r[items.union.fields.intro] Unions have no notion of an "active field". Instead, every union access just interprets the storage as the type of the field used for the access.

r[items.union.fields.read] Reading a union field reads the bits of the union at the field's type.

r[items.union.fields.offset] Fields might have a non-zero offset (except when <u>the C representation</u> is used); in that case the bits starting at the offset of the fields are read

r[items.union.fields.validity] It is the programmer's responsibility to make sure that the data is valid at the field's type. Failing to do so results in <u>undefined behavior</u>. For example, reading the value **3** from a field of the <u>boolean type</u> is undefined behavior. Effectively, writing to and then reading from a union with <u>the C representation</u> is analogous to a <u>transmute</u> from the type used for writing to the type used for reading.

r[items.union.fields.read-safety] Consequently, all reads of union fields have to be placed in unsafe blocks:

```
# union MyUnion { f1: u32, f2: f32 }
# let u = MyUnion { f1: 1 };
#
unsafe {
    let f = u.f1;
}
```

Commonly, code using unions will provide safe wrappers around unsafe union field accesses.

r[items.union.fields.write-safety] In contrast, writes to union fields are safe, since they just overwrite arbitrary data, but cannot cause undefined behavior. (Note that union field types can never have drop glue, so a union field write will never implicitly drop anything.)

r[items.union.pattern]

### Pattern matching on unions

r[items.union.pattern.intro] Another way to access union fields is to use pattern matching.

r[items.union.pattern.one-field] Pattern matching on union fields uses the same syntax as struct patterns, except that the pattern must specify exactly one field.

r[items.union.pattern.safety] Since pattern matching is like reading the union with a particular field, it has to be placed in unsafe blocks as well.

```
# union MyUnion { f1: u32, f2: f32 }
#
fn f(u: MyUnion) {
    unsafe {
        match u {
            MyUnion { f1: 10 } => { println!("ten"); }
            MyUnion { f2 } => { println!("{}", f2); }
        }
    }
}
```

r[items.union.pattern.subpattern] Pattern matching may match a union as a field of a larger structure. In particular, when using a Rust union to implement a C tagged union via FFI, this allows matching on the tag and the corresponding field simultaneously:

```
#[repr(u32)]
enum Tag { I, F }
#[repr(C)]
union U {
    i: i32,
    f: f32,
}
#[repr(C)]
struct Value {
```

```
tag: Tag,
    u: U,
}
fn is_zero(v: Value) -> bool {
    unsafe {
        match v {
            Value { tag: Tag::I, u: U { i: 0 } } => true,
            Value { tag: Tag::F, u: U { i: 0 } } => true,
            Value { tag: Tag::F, u: U { f: num } } if num ==
0.0 => true,
            _ => false,
            }
        }
}
```

r[items.union.ref]

#### **References to union fields**

r[items.union.ref.intro] Since union fields share common storage, gaining write access to one field of a union can give write access to all its remaining fields.

r[items.union.ref.borrow] Borrow checking rules have to be adjusted to account for this fact. As a result, if one field of a union is borrowed, all its remaining fields are borrowed as well for the same lifetime.

```
# union MyUnion { f1: u32, f2: f32 }
// ERROR: cannot borrow `u` (via `u.f2`) as mutable more than
once at a time
fn test() {
    let mut u = MyUnion \{ f1: 1 \};
    unsafe {
        let b1 = \&mut u.f1;
                          ---- first mutable borrow occurs here
11
(via `u.f1`)
        let b_2 = \&mut u.f_2;
11
                         ^^^^ second mutable borrow occurs here
(via `u.f2`)
        *b1 = 5;
    }
// - first borrow ends here
    assert_eq!(unsafe { u.f1 }, 5);
}
```

r[items.union.ref.usage] As you could see, in many aspects (except for layouts, safety, and ownership) unions behave exactly like structs, largely as a consequence of inheriting their syntactic shape from structs. This is also true for many unmentioned aspects of Rust language (such as privacy, name resolution, type inference, generics, trait implementations, inherent implementations, coherence, pattern checking, etc etc etc). r[items.const]

## **Constant items**

r[items.const.syntax]

```
ConstantItem ->
```

```
`const` ( IDENTIFIER | `_` ) `:` Type ( `=` Expression )?
`;`
```

r[items.const.intro] A *constant item* is an optionally named <u>constant</u> <u>value</u> which is not associated with a specific memory location in the program.

r[items.const.behavior] Constants are essentially inlined wherever they are used, meaning that they are copied directly into the relevant context when used. This includes usage of constants from external crates, and non-Copy types. References to the same constant are not necessarily guaranteed to refer to the same memory address.

r[items.const.namespace] The constant declaration defines the constant value in the <u>value namespace</u> of the module or block where it is located.

r[items.const.static] Constants must be explicitly typed. The type must have a 'static lifetime: any references in the initializer must have 'static lifetimes. References in the type of a constant default to 'static lifetime; see static lifetime elision.

r[items.const.static-temporary] A reference to a constant will have 'static lifetime if the constant value is eligible for <u>promotion</u>; otherwise, a temporary will be created.

```
const BIT1: u32 = 1 << 0;
const BIT2: u32 = 1 << 1;
const BITS: [u32; 2] = [BIT1, BIT2];
const STRING: &'static str = "bitstring";
struct BitsNStrings<'a> {
    mybits: [u32; 2],
    mystring: &'a str,
}
```

```
const BITS_N_STRINGS: BitsNStrings<'static> = BitsNStrings {
    mybits: BITS,
    mystring: STRING,
};
```

r[items.const.no-mut-refs] The final value of a **const** item cannot contain any mutable references.

```
# #![allow(static_mut_refs)]
static mut S: u8 = 0;
const C: &u8 = unsafe { &mut S }; // OK
# use core::sync::atomic::AtomicU8;
static S: AtomicU8 = AtomicU8::new(0);
const C: &AtomicU8 = &S; // OK
# #![allow(static_mut_refs)]
static mut S: u8 = 0;
const C: &mut u8 = unsafe { &mut S }; // ERROR not allowed
```

[!NOTE] We also disallow, in the final value, shared references to mutable statics created in the initializer for a separate reason. Consider:

```
# use core::sync::atomic::AtomicU8;
const C: &AtomicU8 = &AtomicU8::new(0); // ERROR
```

Here, the AtomicU8 is a temporary that is lifetime extended to 'static (see [destructors.scope.lifetime-extension.static]), and references to lifetime-extended temporaries with interior mutability are not allowed in the final value of a constant expression (see [const-eval.const-expr.borrows]).

r[items.const.expr-omission] The constant expression may only be omitted in a <u>trait definition</u>.

r[items.const.destructor]

#### **Constants with Destructors**

Constants can contain destructors. Destructors are run when the value goes out of scope.

```
struct TypeWithDestructor(i32);
impl Drop for TypeWithDestructor {
    fn drop(&mut self) {
        println!("Dropped. Held {}.", self.0);
    }
}
const ZERO_WITH_DESTRUCTOR: TypeWithDestructor
TypeWithDestructor(0);
fn create_and_drop_zero_with_destructor() {
    let x = ZERO_WITH_DESTRUCTOR;
    // x gets dropped at end of function, calling drop.
    // prints "Dropped. Held 0.".
}
```

=

r[items.const.unnamed]

#### **Unnamed constant**

r[items.const.unnamed.intro] Unlike an <u>associated constant</u>, a <u>free</u> constant may be unnamed by using an underscore instead of the name. For example:

```
const _: () = { struct _SameNameTwice; };
// OK although it is the same name as above:
const _: () = { struct _SameNameTwice; };
```

r[items.const.unnamed.repetition] As with <u>underscore imports</u>, macros may safely emit the same unnamed constant in the same scope more than once. For example, the following should not produce an error:

```
macro_rules! m {
   ($item: item) => { $item $item }
}
m!(const _: () = (););
// This expands to:
// const _: () = ();
// const _: () = ();
r[items.const.eval]
```

#### **Evaluation**

<u>Free</u> constants are always <u>evaluated</u> at compile-time to surface panics. This happens even within an unused function:

```
// Compile-time panic
const PANIC: () = std::unimplemented!();
```

```
fn unused_generic_function<T>() {
    // A failing compile-time assertion
    const _: () = assert!(usize::BITS == 0);
}
```

r[items.static]

## **Static items**

r[items.static.syntax]

```
StaticItem ->
    ItemSafety?[^extern-safety] `static` `mut`? IDENTIFIER `:`
Type ( `=` Expression )? `;`
1
```

The safe and unsafe function qualifiers are only allowed semantically within extern blocks.

r[items.static.intro] A *static item* is similar to a <u>constant</u>, except that it represents an allocated object in the program that is initialized with the initializer expression. All references and raw pointers to the static refer to the same allocated object.

r[items.static.lifetime] Static items have the static lifetime, which outlives all other lifetimes in a Rust program. Static items do not call <u>drop</u> at the end of the program.

r[items.static.storage-disjointness] If the static has a size of at least 1 byte, this allocated object is disjoint from all other such static objects as well as heap allocations and stack-allocated variables. However, the storage of immutable static items can overlap with objects that do not themselves have a unique address, such as promoteds and const\_items.

r[items.static.namespace] The static declaration defines a static value in the <u>value namespace</u> of the module or block where it is located.

r[items.static.init] The static initializer is a <u>constant expression</u> evaluated at compile time. Static initializers may refer to and read from other statics. When reading from mutable statics, they read the initial value of that static.

r[items.static.read-only] Non-mut static items that contain a type that is not <u>interior mutable</u> may be placed in read-only memory.

r[items.static.safety] All access to a static is safe, but there are a number of restrictions on statics:

r[items.static.sync]
• The type must have the <u>Sync</u> trait bound to allow thread-safe access.

r[items.static.init.omission] The initializer expression must be omitted in an <u>external block</u>, and must be provided for free static items.

r[items.static.safety-qualifiers] The safe and unsafe qualifiers are semantically only allowed when used in an <u>external block</u>.

r[items.static.generics]

#### **Statics & generics**

A static item defined in a generic scope (for example in a blanket or default implementation) will result in exactly one static item being defined, as if the static definition was pulled out of the current scope into the module. There will *not* be one item per monomorphization.

```
This code:
use std::sync::atomic::{AtomicUsize, Ordering};
trait Tr {
    fn default_impl() {
        static COUNTER: AtomicUsize = AtomicUsize::new(0);
                  println!("default_impl:
                                           counter
                                                      was
                                                            {}",
COUNTER.fetch_add(1, Ordering::Relaxed));
    }
    fn blanket_impl();
}
struct Ty1 {}
struct Ty2 {}
impl<T> Tr for T {
    fn blanket_impl() {
        static COUNTER: AtomicUsize = AtomicUsize::new(0);
                  println!("blanket_impl:
                                            counter
                                                      was
                                                            {}",
COUNTER.fetch_add(1, Ordering::Relaxed));
    }
}
fn main() {
    <Ty1 as Tr>::default_impl();
    <Ty2 as Tr>::default_impl();
    <Ty1 as Tr>::blanket_impl();
```

```
<Ty2 as Tr>::blanket_impl();
```

prints

}

```
default_impl: counter was 0
default_impl: counter was 1
blanket_impl: counter was 0
blanket_impl: counter was 1
r[items.static.mut]
```

#### **Mutable statics**

r[items.static.mut.intro] If a static item is declared with the mut keyword, then it is allowed to be modified by the program. One of Rust's goals is to make concurrency bugs hard to run into, and this is obviously a very large source of race conditions or other bugs.

r[items.static.mut.safety] For this reason, an unsafe block is required when either reading or writing a mutable static variable. Care should be taken to ensure that modifications to a mutable static are safe with respect to other threads running in the same process.

r[items.static.mut.extern] Mutable statics are still very useful, however. They can be used with C libraries and can also be bound from C libraries in an extern block.

```
# fn atomic_add(_: *mut u32, _: u32) -> u32 { 2 }
static mut LEVELS: u32 = 0;
// This violates the idea of no shared state, and this doesn't
internally
// protect against races, so this function is `unsafe`
unsafe fn bump_levels_unsafe() -> u32 {
    unsafe {
        let ret = LEVELS;
        LEVELS += 1;
        return ret;
    }
}
// As an alternative to `bump_levels_unsafe`, this function is
safe, assuming
// that we have an atomic_add function which returns the old
value. This
// function is safe only if no other code accesses the static
in a non-atomic
```

```
// fashion. If such accesses are possible (such as in
`bump_levels_unsafe`),
// then this would need to be `unsafe` to indicate to the
caller that they
// must still guard against concurrent access.
fn bump_levels_safe() -> u32 {
    unsafe {
        return atomic_add(&raw mut LEVELS, 1);
    }
}
```

r[items.static.mut.sync] Mutable statics have the same restrictions as normal statics, except that the type does not have to implement the Sync trait.

r[items.static.alternate]

# **Using Statics or Consts**

It can be confusing whether or not you should use a constant item or a static item. Constants should, in general, be preferred over statics unless one of the following are true:

- Large amounts of data are being stored.
- The single-address property of statics is required.
- Interior mutability is required.

r[items.traits]

# Traits

r[items.traits.syntax]

Trait ->

`unsafe`? `trait` IDENTIFIER GenericParams? ( `:` TypeParamBounds? )? WhereClause?

```
`{`
InnerAttribute*
AssociatedItem*
```

r[items.traits.intro] A *trait* describes an abstract interface that types can implement. This interface consists of <u>associated items</u>, which come in three varieties:

- <u>functions</u>
- <u>types</u>
- <u>constants</u>

r[items.traits.namespace] The trait declaration defines a trait in the <u>type</u> <u>namespace</u> of the module or block where it is located.

r[items.traits.associated-item-namespaces] Associated items are defined as members of the trait within their respective namespaces. Associated types are defined in the type namespace. Associated constants and associated functions are defined in the value namespace.

r[items.traits.self-param] All traits define an implicit type parameter Self that refers to "the type that is implementing this interface". Traits may also contain additional type parameters. These type parameters, including Self, may be constrained by other traits and so forth <u>as usual</u>.

r[items.traits.impls] Traits are implemented for specific types through separate <u>implementations</u>.

r[items.traits.associated-item-decls] Trait functions may omit the function body by replacing it with a semicolon. This indicates that the implementation must define the function. If the trait function defines a body, this definition acts as a default for any implementation which does

not override it. Similarly, associated constants may omit the equals sign and expression to indicate implementations must define the constant value. Associated types must never define the type, the type may only be specified in an implementation.

```
// Examples of associated trait items with and without
definitions.
trait Example {
    const CONST_NO_DEFAULT: i32;
    const CONST_WITH_DEFAULT: i32 = 99;
    type TypeNoDefault;
    fn method_without_default(&self);
    fn method_with_default(&self) {}
}
```

r[items.traits.const-fn] Trait functions are not allowed to be <u>const</u>. r[items.traits.bounds]

# **Trait bounds**

Generic items may use traits as <u>bounds</u> on their type parameters. r[items.traits.generic]

### **Generic traits**

Type parameters can be specified for a trait to make it generic. These appear after the trait name, using the same syntax used in <u>generic functions</u>.

```
trait Seq<T> {
    fn len(&self) -> u32;
    fn elt_at(&self, n: u32) -> T;
    fn iter<F>(&self, f: F) where F: Fn(T);
}
```

r[items.traits.dyn-compatible]

# Dyn compatibility

r[items.traits.dyn-compatible.intro] A dyn-compatible trait can be the base trait of a <u>trait object</u>. A trait is *dyn compatible* if it has the following qualities:

r[items.traits.dyn-compatible.supertraits]

- All <u>supertraits</u> must also be dyn compatible. r[items.traits.dyn-compatible.sized]
- Sized must not be a <u>supertrait</u>. In other words, it must not require Self: Sized.

r[items.traits.dyn-compatible.associated-consts]

- It must not have any associated constants. r[items.traits.dyn-compatible.associated-types]
- It must not have any associated types with generics. r[items.traits.dyn-compatible.associated-functions]
- All associated functions must either be dispatchable from a trait object or be explicitly non-dispatchable:
  - Dispatchable functions must:
    - Not have any type parameters (although lifetime parameters are allowed).
    - Be a <u>method</u> that does not use Self except in the type of the receiver.
    - Have a receiver with one of the following types:
      - &Self (i.e. &self)
      - &mut Self (i.e &mut self)
      - Box<Self>
      - Rc<Self>

- Arc<Self>
- <u>Pin<P></u> where P is one of the types above
- Not have an opaque return type; that is,
  - Not be an async fn (which has a hidden Future type).
  - Not have a return position impl Trait type (fn example(&self) -> impl Trait).
- Not have a where Self: Sized bound (receiver type of Self (i.e. self) implies this).
- Explicitly non-dispatchable functions require:
  - Have a where Self: Sized bound (receiver type of Self (i.e. self) implies this).

r[items.traits.dyn-compatible.async-traits]

• The [AsyncFn], [AsyncFnMut], and [AsyncFnOnce] traits are not dyncompatible.

[!NOTE] This concept was formerly known as *object safety*.

```
# use std::rc::Rc;
# use std::sync::Arc;
# use std::pin::Pin;
// Examples of dyn compatible methods.
trait TraitMethods {
    fn by_ref(self: &Self) {}
    fn by_ref_mut(self: &mut Self) {}
    fn by_box(self: Box<Self>) {}
    fn by_box(self: Ro<Self>) {}
    fn by_rc(self: Rc<Self>) {}
    fn by_arc(self: Arc<Self>) {}
    fn by_pin(self: Pin<&Self>) {}
    fn with_lifetime<'a>(self: &'a Self) {}
    fn nested_pin(self: Pin<Arc<Self>>) {}
}
# struct S;
```

```
# impl TraitMethods for S {}
 # let t: Box<dyn TraitMethods> = Box::new(S);
// This trait is dyn compatible, but these methods cannot be
dispatched on a trait object.
trait NonDispatchable {
    // Non-methods cannot be dispatched.
    fn foo() where Self: Sized {}
    // Self type isn't known until runtime.
    fn returns(&self) -> Self where Self: Sized;
      // `other` may be a different concrete type of the
receiver.
    fn param(&self, other: Self) where Self: Sized {}
    // Generics are not compatible with vtables.
    fn typed<T>(&self, x: T) where Self: Sized {}
}
struct S;
impl NonDispatchable for S {
    fn returns(&self) -> Self where Self: Sized { S }
}
let obj: Box<dyn NonDispatchable> = Box::new(S);
obj.returns(); // ERROR: cannot call with Self return
obj.param(S); // ERROR: cannot call with Self parameter
obj.typed(1); // ERROR: cannot call with generic type
# use std::rc::Rc;
// Examples of dyn-incompatible traits.
trait DynIncompatible {
     const CONST: i32 = 1; // ERROR: cannot have associated
const
    fn foo() {} // ERROR: associated function without Sized
    fn returns(&self) -> Self; // ERROR: Self in return type
     fn typed<T>(&self, x: T) {} // ERROR: has generic type
parameters
```

fn nested(self: Rc<Box<Self>>) {} // ERROR: nested receiver

```
cannot be downcasted
}
struct S;
impl DynIncompatible for S {
    fn returns(&self) -> Self { S }
}
let obj: Box<dyn DynIncompatible> = Box::new(S); // ERROR
// `Self: Sized` traits are dyn-incompatible.
trait TraitWithSize where Self: Sized {}
struct S;
impl TraitWithSize for S {}
let obj: Box<dyn TraitWithSize> = Box::new(S); // ERROR
// Dyn-incompatible if `Self` is a type argument.
trait Super<A> {}
trait WithSelf: Super<Self> where Self: Sized {}
struct S;
impl<A> Super<A> for S {}
impl WithSelf for S {}
let obj: Box<dyn WithSelf> = Box::new(S); // ERROR: cannot use
`Self` type parameter
  r[items.traits.supertraits]
```

#### **Supertraits**

r[items.traits.supertraits.intro] **Supertraits** are traits that are required to be implemented for a type to implement a specific trait. Furthermore, anywhere a <u>generic</u> or <u>trait object</u> is bounded by a trait, it has access to the associated items of its supertraits.

r[items.traits.supertraits.decl] Supertraits are declared by trait bounds on the Self type of a trait and transitively the supertraits of the traits declared in those trait bounds. It is an error for a trait to be its own supertrait.

r[items.traits.supertraits.subtrait] The trait with a supertrait is called a **subtrait** of its supertrait.

The following is an example of declaring Shape to be a supertrait of Circle.

```
trait Shape { fn area(&self) -> f64; }
trait Circle: Shape { fn radius(&self) -> f64; }
```

And the following is the same example, except using <u>where clauses</u>.

```
trait Shape { fn area(&self) -> f64; }
```

```
trait Circle where Self: Shape { fn radius(&self) -> f64; }
```

This next example gives radius a default implementation using the area function from Shape.

```
# trait Shape { fn area(&self) -> f64; }
trait Circle where Self: Shape {
   fn radius(&self) -> f64 {
        // A = pi * r^2
        // so algebraically,
        // r = sqrt(A / pi)
        (self.area() / std::f64::consts::PI).sqrt()
     }
}
```

}

This next example calls a supertrait method on a generic parameter.

```
# trait Shape { fn area(&self) -> f64; }
# trait Circle: Shape { fn radius(&self) -> f64; }
```

```
fn print_area_and_radius<C: Circle>(c: C) {
    // Here we call the area method from the supertrait
`Shape` of `Circle`.
    println!("Area: {}", c.area());
    println!("Radius: {}", c.radius());
}
```

Similarly, here is an example of calling supertrait methods on trait objects.

```
# trait Shape { fn area(&self) -> f64; }
# trait Circle: Shape { fn radius(&self) -> f64; }
# struct UnitCircle;
# impl Shape for UnitCircle { fn area(&self) -> f64 {
std::f64::consts::PI } }
# impl Circle for UnitCircle { fn radius(&self) -> f64 { 1.0 }
}
# let circle = UnitCircle;
let circle = Box::new(circle) as Box<dyn Circle>;
let nonsense = circle.radius() * circle.area();
```

```
r[items.traits.safety]
```

#### **Unsafe traits**

r[items.traits.safety.intro] Traits items that begin with the unsafe keyword indicate that *implementing* the trait may be <u>unsafe</u>. It is safe to use a correctly implemented unsafe trait. The <u>trait implementation</u> must also begin with the unsafe keyword.

Sync and Send are examples of unsafe traits.

r[items.traits.params]

#### **Parameter patterns**

r[items.traits.params.patterns-no-body] Parameters in associated functions without a body only allow [IDENTIFIER] or <u>wild card</u> patterns, as well as the form allowed by [SelfParam]. mut [IDENTIFIER] is currently allowed, but it is deprecated and will become a hard error in the future.

```
trait T {
    fn f1(&self);
    fn f2(x: Self, _: i32);
}
```

```
trait T {
```

fn f2(&x: &i32); // ERROR: patterns aren't allowed in
functions without bodies

}

r[items.traits.params.patterns-with-body] Parameters in associated functions with a body only allow irrefutable patterns.

```
trait T {
    fn f1((a, b): (i32, i32)) {} // OK: is irrefutable
}
trait T {
    fn f1(123: i32) {} // ERROR: pattern is refutable
        fn f2(Some(x): Option<i32>) {} // ERROR: pattern is
refutable
}
```

r[items.traits.params.pattern-required.edition2018]

```
[!EDITION-2018] Prior to the 2018 edition, the pattern for an associated function parameter is optional:
```

```
// 2015 Edition
trait T {
    fn f(i32); // OK: parameter identifiers are not
required
}
```

Beginning in the 2018 edition, patterns are no longer optional.

r[items.traits.params.restriction-patterns.edition2018]

[!EDITION-2018] Prior to the 2018 edition, parameters in associated functions with a body are limited to the following kinds of patterns:

- [IDENTIFIER]
- mut [IDENTIFIER]
- \_
- & [IDENTIFIER]
- && [IDENTIFIER]

```
// 2015 Edition
```

```
trait T {
```

```
fn f1((a, b): (i32, i32)) {} // ERROR: pattern not
allowed
```

}

Beginning in 2018, all irrefutable patterns are allowed as described in [items.traits.params.patterns-with-body].

r[items.traits.associated-visibility]

# Item visibility

r[items.traits.associated-visibility.intro] Trait items syntactically allow a [Visibility] annotation, but this is rejected when the trait is validated. This allows items to be parsed with a unified syntax across different contexts where they are used. As an example, an empty vis macro fragment specifier can be used for trait items, where the macro rule may be used in other situations where visibility is allowed.

```
macro_rules! create_method {
    ($vis:vis $name:ident) => {
        $vis fn $name(&self) {}
    };
}
trait T1 {
    // Empty `vis` is allowed.
    create_method! { method_of_t1 }
}
struct S;
impl S {
    // Visibility is allowed here.
    create_method! { pub method_of_s }
}
impl T1 for S {}
fn main() {
    let s = S;
    s.method_of_t1();
    s.method_of_s();
}
```

r[items.impl]

# Implementations

```
r[items.impl.syntax]
Implementation -> InherentImpl | TraitImpl
InherentImpl ->
    `impl` GenericParams? Type WhereClause? `{`
        InnerAttribute*
        AssociatedItem*
    `}`
TraitImpl ->
    `unsafe`? `impl` GenericParams? `!`? TypePath `for` Type
    WhereClause?
    `{`
        InnerAttribute*
        AssociatedItem*
    `}`
```

r[items.impl.intro] An *implementation* is an item that associates items with an *implementing type*. Implementations are defined with the keyword impl and contain functions that belong to an instance of the type that is being implemented or to the type statically.

r[items.impl.kinds] There are two types of implementations:

- inherent implementations
- <u>trait</u> implementations

r[items.impl.inherent]

# **Inherent Implementations**

r[items.impl.inherent.intro] An inherent implementation is defined as the sequence of the impl keyword, generic type declarations, a path to a nominal type, a where clause, and a bracketed set of associable items.

r[items.impl.inherent.implementing-type] The nominal type is called the *implementing type* and the associable items are the *associated items* to the implementing type.

r[items.impl.inherent.associated-items] Inherent implementations associate the contained items to the implementing type.

r[items.impl.inherent.associated-items.allowed-items] Inherent implementations can contain <u>associated functions</u> (including <u>methods</u>) and <u>associated constants</u>.

r[items.impl.inherent.type-alias] They cannot contain associated type aliases.

r[items.impl.inherent.associated-item-path] The <u>path</u> to an associated item is any path to the implementing type, followed by the associated item's identifier as the final path component.

r[items.impl.inherent.coherence] A type can also have multiple inherent implementations. An implementing type must be defined within the same crate as the original type definition.

```
pub mod color {
   pub struct Color(pub u8, pub u8, pub u8);
   impl Color {
      pub const WHITE: Color = Color(255, 255, 255);
   }
}
mod values {
   use super::color::Color;
   impl Color {
      pub fn red() -> Color {
        Color(255, 0, 0)
   }
}
```

```
}
   }
}
pub use self::color::Color;
fn main() {
    // Actual path to the implementing type and impl in the
same module.
   color::Color::WHITE;
     // Impl blocks in different modules are still accessed
through a path to the type.
    color::Color::red();
   // Re-exported paths to the implementing type also work.
   Color::red();
   // Does not work, because use in `values` is not pub.
   // values::Color::red();
}
```

r[items.impl.trait]

## **Trait Implementations**

r[items.impl.trait.intro] A *trait implementation* is defined like an inherent implementation except that the optional generic type declarations are followed by a <u>trait</u>, followed by the keyword <code>for</code>, followed by a path to a nominal type.

r[items.impl.trait.implemented-trait] The trait is known as the *implemented trait*. The implementing type implements the implemented trait.

r[items.impl.trait.def-requirement] A trait implementation must define all non-default associated items declared by the implemented trait, may redefine default associated items defined by the implemented trait, and cannot define any other items.

r[items.impl.trait.associated-item-path] The path to the associated items is < followed by a path to the implementing type followed by as followed by a path to the trait followed by > as a path component followed by the associated item's path component.

r[items.impl.trait.safety] <u>Unsafe traits</u> require the trait implementation to begin with the unsafe keyword.

```
# #[derive(Copy, Clone)]
# struct Point {x: f64, y: f64};
# type Surface = i32;
# struct BoundingBox {x: f64, y: f64, width: f64, height:
f64};
    trait
                    {
                        fn
                             draw(&self,
                                                Surface);
            Shape
                                                             fn
#
                                           s:
bounding_box(&self) -> BoundingBox; }
# fn do_draw_circle(s: Surface, c: Circle) { }
struct Circle {
    radius: f64,
    center: Point,
}
impl Copy for Circle {}
```

```
impl Clone for Circle {
    fn clone(&self) -> Circle { *self }
}
impl Shape for Circle {
    fn draw(&self, s: Surface) { do_draw_circle(s, *self); }
    fn bounding_box(&self) -> BoundingBox {
        let r = self.radius;
        BoundingBox {
            x: self.center.x - r,
            y: self.center.y - r,
            width: 2.0 * r,
            height: 2.0 * r,
        }
    }
}
```

r[items.impl.trait.coherence]

### **Trait Implementation Coherence**

r[items.impl.trait.coherence.intro] A trait implementation is considered incoherent if either the orphan rules check fails or there are overlapping implementation instances.

r[items.impl.trait.coherence.overlapping] Two trait implementations overlap when there is a non-empty intersection of the traits the implementation is for, the implementations can be instantiated with the same type.

r[items.impl.trait.orphan-rule]

#### **Orphan rules**

r[items.impl.trait.orphan-rule.intro] The *orphan rule* states that a trait implementation is only allowed if either the trait or at least one of the types in the implementation is defined in the current crate. It prevents conflicting trait implementations across different crates and is key to ensuring coherence. An orphan implementation is one that implements a foreign trait for a foreign type. If these were freely allowed, two crates could implement the same trait for the same type in incompatible ways, creating a situation where adding or updating a dependency could break compilation due to conflicting implementations.

The orphan rule enables library authors to add new implementations to their traits without fear that they'll break downstream code. Without these restrictions, a library couldn't add an implementation like impl<T: Display> MyTrait for T without potentially conflicting with downstream implementations.

r[items.impl.trait.orphan-rule.general] Given impl<P1..=Pn> Trait<T1..=Tn> for T0, an impl is valid only if at least one of the following is true:

- Trait is a <u>local trait</u>
- All of
  - At least one of the types T0..=Tn must be a <u>local type</u>. Let Ti be the first such type.
  - No <u>uncovered type</u> parameters P1..=Pn may appear in T0..Ti (excluding Ti)

r[items.impl.trait.uncovered-param] Only the appearance of *uncovered* type parameters is restricted.

r[items.impl.trait.fundamental] Note that for the purposes of coherence, <u>fundamental types</u> are special. The T in Box<T> is not considered covered, and Box<LocalType> is considered local.

r[items.impl.generics]

#### **Generic Implementations**

r[items.impl.generics.intro] An implementation can take <u>generic</u> <u>parameters</u>, which can be used in the rest of the implementation. Implementation parameters are written directly after the <u>impl</u> keyword.

```
# trait Seq<T> { fn dummy(&self, _: T) { } }
impl<T> Seq<T> for Vec<T> {
    /* ... */
}
impl Seq<bool> for u32 {
    /* Treat the integer as a sequence of bits */
}
```

r[items.impl.generics.usage] Generic parameters *constrain* an implementation if the parameter appears at least once in one of:

- The implemented trait, if it has one
- The implementing type
- As an <u>associated type</u> in the <u>bounds</u> of a type that contains another parameter that constrains the implementation

r[items.impl.generics.constrain] Type and const parameters must always constrain the implementation. Lifetimes must constrain the implementation if the lifetime is used in an associated type.

Examples of constraining situations:

```
# trait Trait{}
# trait GenericTrait<T> {}
# trait GenericTrait<T> {}
# trait HasAssocType { type Ty; }
# struct Struct;
# struct GenericStruct<T>(T);
# struct ConstGenericStruct<const N: usize>([(); N]);
// T constrains by being an argument to GenericTrait.
impl<T> GenericTrait<T> for i32 { /* ... */ }
// T constrains by being an argument to GenericStruct
```

```
impl<T> Trait for GenericStruct<T> { /* ... */ }
```

```
//
    Likewise,
               N constrains by being
                                             an
                                                  argument to
ConstGenericStruct
impl<const N: usize> Trait for ConstGenericStruct<N> { /* ...
*/ }
// T constrains by being in an associated type in a bound for
type `U` which is
// itself a generic parameter constraining the trait.
impl<T, U> GenericTrait<U> for u32 where U: HasAssocType<Ty =</pre>
T> { /* ... */ }
// Like previous, except the type is `(U, isize)`. `U` appears
inside the type
// that includes `T`, and is not the type itself.
impl<T, U> GenericStruct<U> where (U, isize): HasAssocType<Ty</pre>
= T > \{ /* ... * / \}
```

Examples of non-constraining situations:

// The rest of these are errors, since they have type or const
parameters that

// do not constrain.

```
// T does not constrain since it does not appear at all.
impl<T> Struct { /* ... */ }
```

```
// N does not constrain for the same reason.
impl<const N: usize> Struct { /* ... */ }
```

```
// Usage of T inside the implementation does not constrain the
impl.
impl<T> Struct {
   fn uses_t(t: &T) { /* ... */ }
}
```

// T is used as an associated type in the bounds for U, but U

```
does not constrain.
impl<T, U> Struct where U: HasAssocType<Ty = T> { /* ... */ }
// T is used in the bounds, but not as an associated type, so
it does not constrain.
```

impl<T, U> GenericTrait<U> for u32 where U: GenericTrait<T> {}

Example of an allowed unconstraining lifetime parameter:

```
# struct Struct;
impl<'a> Struct {}
```

Example of a disallowed unconstraining lifetime parameter:

```
# struct Struct;
# trait HasAssocType { type Ty; }
impl<'a> HasAssocType for Struct {
   type Ty = &'a Struct;
}
```

```
r[items.impl.attributes]
```

# **Attributes on Implementations**

Implementations may contain outer <u>attributes</u> before the <u>impl</u> keyword and inner <u>attributes</u> inside the brackets that contain the associated items. Inner attributes must come before any associated items. The attributes that have meaning here are <u>cfg</u>, <u>deprecated</u>, <u>doc</u>, and <u>the lint check</u> <u>attributes</u>. r[items.extern]

# **External blocks**

```
r[items.extern.syntax]
ExternBlock ->
  `unsafe`?[^unsafe-2024] `extern` Abi? `{`
    InnerAttribute*
    ExternalItem*
    `}`
ExternalItem ->
    OuterAttribute* (
        MacroInvocationSemi
        | Visibility? StaticItem
        | Visibility? Function
    )
```

1

Starting with the 2024 Edition, the unsafe keyword is required semantically.

r[items.extern.intro] External blocks provide *declarations* of items that are not *defined* in the current crate and are the basis of Rust's foreign function interface. These are akin to unchecked imports.

r[items.extern.allowed-kinds] Two kinds of item *declarations* are allowed in external blocks: <u>functions</u> and <u>statics</u>.

r[items.extern.fn-safety] Calling functions or accessing statics that are declared in external blocks is only allowed in an unsafe context.

r[items.extern.namespace] The external block defines its functions and statics in the <u>value namespace</u> of the module or block where it is located.

r[items.extern.unsafe-required] The unsafe keyword is semantically required to appear before the extern keyword on external blocks.

r[items.extern.edition2024]

[!EDITION-2024] Prior to the 2024 edition, the unsafe keyword is optional. The safe and unsafe item qualifiers are only allowed if the

external block itself is marked as unsafe.

r[items.extern.fn]

#### **Functions**

r[items.extern.fn.body] Functions within external blocks are declared in the same way as other Rust functions, with the exception that they must not have a body and are instead terminated by a semicolon.

r[items.extern.fn.param-patterns] Patterns are not allowed in parameters, only [IDENTIFIER] or \_\_\_\_\_ may be used.

r[items.extern.fn.qualifiers] The safe and unsafe function qualifiers are allowed, but other function qualifiers (e.g. const, async, extern) are not.

r[items.extern.fn.foreign-abi] Functions within external blocks may be called by Rust code, just like functions defined in Rust. The Rust compiler automatically translates between the Rust ABI and the foreign ABI.

r[items.extern.fn.safety] A function declared in an extern block is implicitly unsafe unless the safe function qualifier is present.

r[items.extern.fn.fn-ptr] When coerced to a function pointer, a function declared in an extern block has type extern "abi" for<'l1, ..., 'lm> fn(A1, ..., An) -> R, where 'l1, ... 'lm are its lifetime parameters, A1, ..., An are the declared types of its parameters, R is the declared return type.

r[items.extern.static]
#### **Statics**

r[items.extern.static.intro] Statics within external blocks are declared in the same way as <u>statics</u> outside of external blocks, except that they do not have an expression initializing their value.

r[items.extern.static.safety] Unless a static item declared in an extern block is qualified as safe, it is unsafe to access that item, whether or not it's mutable, because there is nothing guaranteeing that the bit pattern at the static's memory is valid for the type it is declared with, since some arbitrary (e.g. C) code is in charge of initializing the static.

r[items.extern.static.mut] Extern statics can be either immutable or mutable just like <u>statics</u> outside of external blocks.

r[items.extern.static.read-only] An immutable static *must* be initialized before any Rust code is executed. It is not enough for the static to be initialized before Rust code reads from it. Once Rust code runs, mutating an immutable static (from inside or outside Rust) is UB, except if the mutation happens to bytes inside of an UnsafeCell.

r[items.extern.abi]

#### ABI

r[items.extern.abi.intro] By default external blocks assume that the library they are calling uses the standard C ABI on the specific platform. Other ABIs may be specified using an abi string, as shown here:

```
// Interface to the Windows API
unsafe extern "system" { }
```

r[items.extern.abi.standard] The following ABI strings are supported on all platforms:

r[items.extern.abi.rust]

• unsafe extern "Rust" -- The default ABI when you write a normal fn foo() in any Rust code.

r[items.extern.abi.c]

- unsafe extern "C" -- This is the same as extern fn foo(); whatever the default your C compiler supports.
   r[items.extern.abi.system]
- unsafe extern "system" -- Usually the same as extern "C", except on Win32, in which case it's "stdcall", or what you should use to link to the Windows API itself

r[items.extern.abi.unwind]

extern "C-unwind" and extern "system-unwind" -- identical to
 "C" and "system", respectively, but with <u>different behavior</u> when the callee unwinds (by panicking or throwing a C++ style exception).

r[items.extern.abi.platform] There are also some platform-specific ABI strings:

r[items.extern.abi.cdecl]

• unsafe extern "cdecl" -- The default for x86\_32 C code. r[items.extern.abi.stdcall] • unsafe extern "stdcall" -- The default for the Win32 API on x86\_32.

r[items.extern.abi.win64]

• unsafe extern "win64" -- The default for C code on x86\_64 Windows.

r[items.extern.abi.sysv64]

• unsafe extern "sysv64" -- The default for C code on non-Windows x86\_64.

r[items.extern.abi.aapcs]

• unsafe extern "aapcs" -- The default for ARM.

r[items.extern.abi.fastcall]

 unsafe extern "fastcall" -- The fastcall ABI -- corresponds to MSVC's \_\_fastcall and GCC and clang's \_\_attribute\_\_((fastcall))

r[items.extern.abi.thiscall]

 unsafe extern "thiscall" -- The default for C++ member functions on x86\_32 MSVC -- corresponds to MSVC's \_\_thiscall and GCC and clang's \_\_attribute\_\_((thiscall))

r[items.extern.abi.efiapi]

• unsafe extern "efiapi" -- The ABI used for <u>UEFI</u> functions.

r[items.extern.abi.platform-unwind-variants] Like "C" and "system", most platform-specific ABI strings also have a <u>corresponding</u>\_<u>-unwind</u> <u>variant</u>; specifically, these are:

- "aapcs-unwind"
- "cdecl-unwind"
- "fastcall-unwind"

- "stdcall-unwind"
- "sysv64-unwind"
- "thiscall-unwind"
- "win64-unwind"

r[items.extern.variadic]

# Variadic functions

Functions within external blocks may be variadic by specifying ... as the last argument. The variadic parameter may optionally be specified with an identifier.

```
unsafe extern "C" {
    unsafe fn foo(...);
    unsafe fn bar(x: i32, ...);
    unsafe fn with_name(format: *const u8, args: ...);
    // SAFETY: This function guarantees it will not access
    // variadic arguments.
    safe fn ignores_variadic_arguments(x: i32, ...);
}
```

[!WARNING] The safe qualifier should not be used on a function in an extern block unless that function guarantees that it will not access the variadic arguments at all. Passing an unexpected number of arguments or arguments of unexpected type to a variadic function may lead to [undefined behavior][undefined].

r[items.extern.attributes]

## **Attributes on extern blocks**

r[items.extern.attributes.intro] The following <u>attributes</u> control the behavior of external blocks.

r[items.extern.attributes.link]

#### The link attribute

r[items.extern.attributes.link.intro] The *link* attribute specifies the name of a native library that the compiler should link with for the items within an extern block.

r[items.extern.attributes.link.syntax] It uses the [MetaListNameValueStr] syntax to specify its inputs. The name key is the name of the native library to link. The kind key is an optional value which specifies the kind of library with the following possible values:

r[items.extern.attributes.link.dylib]

• dylib --- Indicates a dynamic library. This is the default if kind is not specified.

r[items.extern.attributes.link.static]

• static --- Indicates a static library.

r[items.extern.attributes.link.framework]

• framework --- Indicates a macOS framework. This is only valid for macOS targets.

r[items.extern.attributes.link.raw-dylib]

raw-dylib --- Indicates a dynamic library where the compiler will generate an import library to link against (see <u>dylib versus raw-</u><u>dylib</u> below for details). This is only valid for Windows targets.

r[items.extern.attributes.link.name-requirement] The name key must be included if kind is specified.

r[items.extern.attributes.link.modifiers] The optional modifiers argument is a way to specify linking modifiers for the library to link.

r[items.extern.attributes.link.modifiers.syntax] Modifiers are specified as a comma-delimited string with each modifier prefixed with either a + or - to indicate that the modifier is enabled or disabled, respectively.

r[items.extern.attributes.link.modifiers.multiple] Specifying multiple modifiers arguments in a single link attribute, or multiple identical modifiers in the same modifiers argument is not currently supported.

```
Example: #[link(name = "mylib", kind = "static", modifiers =
"+whole-archive")].
```

```
r[items.extern.attributes.link.wasm_import_module] The
wasm_import_module key may be used to specify the <u>WebAssembly</u>
module name for the items within an extern block when importing
symbols from the host environment. The default module name is env if
wasm_import_module is not specified.
```

```
#[link(name = "crypto")]
unsafe extern {
    // ...
}
#[link(name = "CoreFoundation", kind = "framework")]
unsafe extern {
    // ...
}
#[link(wasm_import_module = "foo")]
unsafe extern {
    // ...
}
```

r[items.extern.attributes.link.empty-block] It is valid to add the link attribute on an empty extern block. You can use this to satisfy the linking requirements of extern blocks elsewhere in your code (including upstream crates) instead of adding the attribute to each extern block.

r[items.extern.attributes.link.modifiers.bundle]

#### Linking modifiers: bundle

r[items.extern.attributes.link.modifiers.bundle.allowed-kinds] This modifier is only compatible with the static linking kind. Using any other kind will result in a compiler error.

r[items.extern.attributes.link.modifiers.bundle.behavior] When building a rlib or staticlib +bundle means that the native static library will be packed into the rlib or staticlib archive, and then retrieved from there during linking of the final binary.

r[items.extern.attributes.link.modifiers.bundle.behavior-negative] When building a rlib -bundle means that the native static library is registered as a dependency of that rlib "by name", and object files from it are included only during linking of the final binary, the file search by that name is also performed during final linking.

When building a staticlib -bundle means that the native static library is simply not included into the archive and some higher level build system will need to add it later during linking of the final binary.

r[items.extern.attributes.link.modifiers.bundle.no-effect] This modifier has no effect when building other targets like executables or dynamic libraries.

r[items.extern.attributes.link.modifiers.bundle.default] The default for this modifier is +bundle.

More implementation details about this modifier can be found in <u>bundle documentation for rustc</u>.

r[items.extern.attributes.link.modifiers.whole-archive]

#### Linking modifiers: whole-archive

r[items.extern.attributes.link.modifiers.whole-archive.allowed-kinds] This modifier is only compatible with the static linking kind. Using any other kind will result in a compiler error.

r[items.extern.attributes.link.modifiers.whole-archive.behavior] +wholearchive means that the static library is linked as a whole archive without throwing any object files away. r[items.extern.attributes.link.modifiers.whole-archive.default] The default for this modifier is -whole-archive.

More implementation details about this modifier can be found in <u>whole-</u> <u>archive</u> documentation for rustc.

r[items.extern.attributes.link.modifiers.verbatim]

#### Linking modifiers: verbatim

r[items.extern.attributes.link.modifiers.verbatim.allowed-kinds] This modifier is compatible with all linking kinds.

r[items.extern.attributes.link.modifiers.verbatim.behavior] +verbatim
means that rustc itself won't add any target-specified library prefixes or
suffixes (like lib or .a) to the library name, and will try its best to ask for
the same thing from the linker.

r[items.extern.attributes.link.modifiers.verbatim.behavior-negative] verbatim means that rustc will either add a target-specific prefix and suffix
to the library name before passing it to linker, or won't prevent linker from
implicitly adding it.

r[items.extern.attributes.link.modifiers.verbatim.default] The default for this modifier is -verbatim.

More implementation details about this modifier can be found in <u>verbatim documentation for rustc</u>.

r[items.extern.attributes.link.kind-raw-dylib]

#### dylib versus raw-dylib

r[items.extern.attributes.link.kind-raw-dylib.intro] On Windows, linking against a dynamic library requires that an import library is provided to the linker: this is a special static library that declares all of the symbols exported by the dynamic library in such a way that the linker knows that they have to be dynamically loaded at runtime.

r[items.extern.attributes.link.kind-raw-dylib.import] Specifying kind =
"dylib" instructs the Rust compiler to link an import library based on the
name key. The linker will then use its normal library resolution logic to find
that import library. Alternatively, specifying kind = "raw-dylib" instructs

the compiler to generate an import library during compilation and provide that to the linker instead.

r[items.extern.attributes.link.kind-raw-dylib.platform-specific] rawdylib is only supported on Windows. Using it when targeting other platforms will result in a compiler error.

r[items.extern.attributes.link.import\_name\_type]

#### The import\_name\_type key

r[items.extern.attributes.link.import\_name\_type.intro] On x86 Windows, names of functions are "decorated" (i.e., have a specific prefix and/or suffix added) to indicate their calling convention. For example, a stdcall calling convention function with the name fn1 that has no arguments would be decorated as \_fn1@0. However, the PE Format does also permit names to have no prefix or be undecorated. Additionally, the MSVC and GNU toolchains use different decorations for the same calling conventions which means, by default, some Win32 functions cannot be called using the raw-dylib link kind via the GNU toolchain.

r[items.extern.attributes.link.import\_name\_type.values] To allow for these differences, when using the raw-dylib link kind you may also specify the import\_name\_type key with one of the following values to change how functions are named in the generated import library:

- decorated : The function name will be fully-decorated using the MSVC toolchain format.
- noprefix: The function name will be decorated using the MSVC toolchain format, but skipping the leading ?, @, or optionally \_.
- undecorated: The function name will not be decorated.

r[items.extern.attributes.link.import\_name\_type.default] If the import\_name\_type key is not specified, then the function name will be fully-decorated using the target toolchain's format.

r[items.extern.attributes.link.import\_name\_type.variables] Variables are never decorated and so the import\_name\_type key has no effect on how they are named in the generated import library. r[items.extern.attributes.link.import\_name\_type.platform-specific] The import\_name\_type key is only supported on x86 Windows. Using it when targeting other platforms will result in a compiler error.

r[items.extern.attributes.link\_name]

#### The link\_name attribute

r[items.extern.attributes.link\_name.intro] The *link\_name attribute* may be specified on declarations inside an extern block to indicate the symbol to import for the given function or static.

r[items.extern.attributes.link\_name.syntax] It uses the [MetaNameValueStr] syntax to specify the name of the symbol.

```
unsafe extern {
    #[link_name = "actual_symbol_name"]
    safe fn name_in_rust();
}
```

r[items.extern.attributes.link\_name.exclusive] Using this attribute with the link\_ordinal attribute will result in a compiler error.

r[items.extern.attributes.link\_ordinal]

#### The link\_ordinal attribute

r[items.extern.attributes.link\_ordinal.intro] The *link\_ordinal* attribute can be applied on declarations inside an extern block to indicate the numeric ordinal to use when generating the import library to link against. An ordinal is a unique number per symbol exported by a dynamic library on Windows and can be used when the library is being loaded to find that symbol rather than having to look it up by name.

[!WARNING] link\_ordinal should only be used in cases where the ordinal of the symbol is known to be stable: if the ordinal of a symbol is not explicitly set when its containing binary is built then one will be automatically assigned to it, and that assigned ordinal may change between builds of the binary.

```
# #[cfg(all(windows, target_arch = "x86"))]
#[link(name = "exporter", kind = "raw-dylib")]
```

```
unsafe extern "stdcall" {
    #[link_ordinal(15)]
    safe fn imported_function_stdcall(i: i32);
}
```

r[items.extern.attributes.link\_ordinal.allowed-kinds] This attribute is only used with the raw-dylib linking kind. Using any other kind will result in a compiler error.

r[items.extern.attributes.link\_ordinal.exclusive] Using this attribute with the link\_name attribute will result in a compiler error.

r[items.extern.attributes.fn-parameters]

#### **Attributes on function parameters**

Attributes on extern function parameters follow the same rules and restrictions as <u>regular function parameters</u>.

r[items.generics]

# **Generic parameters**

r[items.generics.syntax]

```
GenericParams -> `<` ( GenericParam (`,` GenericParam)* `,`? )?
`>`
```

GenericParam -> OuterAttribute\* ( LifetimeParam | TypeParam | ConstParam )

LifetimeParam -> Lifetime ( `:` LifetimeBounds )?

```
TypeParam -> IDENTIFIER ( `:` TypeParamBounds? )? ( `=` Type )?
```

ConstParam ->

`const` IDENTIFIER `:` Type

( `=` BlockExpression | IDENTIFIER | `-`?LiteralExpression
)?

r[items.generics.syntax.intro] <u>Functions</u>, <u>type aliases</u>, <u>structs</u>, <u>enumerations</u>, <u>unions</u>, <u>traits</u>, and <u>implementations</u> may be *parameterized* by types, constants, and lifetimes. These parameters are listed in angle brackets (<...>), usually immediately after the name of the item and before its definition. For implementations, which don't have a name, they come directly after impl.

r[items.generics.syntax.decl-order] The order of generic parameters is restricted to lifetime parameters and then type and const parameters intermixed.

r[items.generics.syntax.duplicate-params] The same parameter name may not be declared more than once in a [GenericParams] list.

Some examples of items with type, const, and lifetime parameters:

```
fn foo<'a, T>() {}
trait A<U> {}
struct Ref<'a, T> where T: 'a { r: &'a T }
```

```
struct InnerArray<T, const N: usize>([T; N]);
struct EitherOrderWorks<const N: bool, U>(U);
```

r[items.generics.syntax.scope] Generic parameters are in scope within the item definition where they are declared. They are not in scope for items declared within the body of a function as described in <u>item declarations</u>. See <u>generic parameter scopes</u> for more details.

r[items.generics.builtin-generic-types] <u>References</u>, <u>raw pointers</u>, <u>arrays</u>, <u>slices</u>, <u>tuples</u>, and <u>function pointers</u> have lifetime or type parameters as well, but are not referred to with path syntax.

r[items.generics.invalid-lifetimes] '\_ and 'static are not valid lifetime parameter names.

r[items.generics.const]

#### **Const generics**

r[items.generics.const.intro] *Const generic parameters* allow items to be generic over constant values.

r[items.generics.const.namespace] The const identifier introduces a name in the <u>value namespace</u> for the constant parameter, and all instances of the item must be instantiated with a value of the given type.

r[items.generics.const.allowed-types] The only allowed types of const parameters are u8, u16, u32, u64, u128, usize, i8, i16, i32, i64, i128, isize, char and bool.

r[items.generics.const.usage] Const parameters can be used anywhere a <u>const item</u> can be used, with the exception that when used in a <u>type</u> or <u>array</u> <u>repeat expression</u>, it must be standalone (as described below). That is, they are allowed in the following places:

- 1. As an applied const to any type which forms a part of the signature of the item in question.
- 2. As part of a const expression used to define an <u>associated const</u>, or as a parameter to an <u>associated type</u>.
- 3. As a value in any runtime expression in the body of any functions in the item.

4. As a parameter to any type used in the body of any functions in the item.

```
5. As a part of the type of any fields in the item.
```

```
// Examples where const generic parameters can be used.
 // Used in the signature of the item itself.
 fn foo<const N: usize>(arr: [i32; N]) {
     // Used as a type within a function body.
     let x: [i32; N];
     // Used as an expression.
     println!("{}", N * 2);
 }
 // Used as a field of a struct.
 struct Foo<const N: usize>([i32; N]);
 impl<const N: usize> Foo<N> {
     // Used as an associated constant.
     const CONST: usize = N * 4;
 }
 trait Trait {
     type Output;
 }
 impl<const N: usize> Trait for Foo<N> {
     // Used as an associated type.
     type Output = [i32; N];
 }
// Examples where const generic parameters cannot be used.
fn foo<const N: usize>() {
   // Cannot use in item definitions within a function body.
    const BAD_CONST: [usize; N] = [1; N];
    static BAD_STATIC: [usize; N] = [1; N];
    fn inner(bad_arg: [usize; N]) {
```

```
let bad_value = N * 2;
}
type BadAlias = [usize; N];
struct BadStruct([usize; N]);
```

}

r[items.generics.const.standalone] As a further restriction, const parameters may only appear as a standalone argument inside of a <u>type</u> or <u>array repeat expression</u>. In those contexts, they may only be used as a single segment <u>path expression</u>, possibly inside a <u>block</u> (such as N or {N}). That is, they cannot be combined with other expressions.

// Examples where const parameters may not be used.

```
// Not allowed to combine in other expressions in types, such
as the
// arithmetic expression in the return type here.
fn bad_function<const N: usize>() -> [u8; {N + 1}] {
    // Similarly not allowed for array repeat expressions.
    [1; {N + 1}]
}
```

r[items.generics.const.argument] A const argument in a <u>path</u> specifies the const value to use for that item.

r[items.generics.const.argument.const-expr] The argument must either be an <u>inferred const</u> or be a <u>const expression</u> of the type ascribed to the const parameter. The const expression must be a <u>block expression</u> (surrounded with braces) unless it is a single path segment (an [IDENTIFIER]) or a <u>literal</u> (with a possibly leading - token).

[!NOTE] This syntactic restriction is necessary to avoid requiring infinite lookahead when parsing an expression inside of a type.

```
struct S<const N: i64>;
const C: i64 = 1;
fn f<const N: i64>() -> S<N> { S }
let _ = f::<1>(); // Literal.
let _ = f::<-1>(); // Negative literal.
```

```
let _ = f::<{ 1 + 2 }>(); // Constant expression.
let _ = f::<C>(); // Single segment path.
let _ = f::<{ C + 1 }>(); // Constant expression.
let _: S<1> = f::<_>(); // Inferred const.
let _: S<1> = f::<(((_)))>(); // Inferred const.
```

[!NOTE] In a generic argument list, an <u>inferred const</u> is parsed as an [inferred type][InferredType] but then semantically treated as a separate kind of <u>const generic argument</u>.

r[items.generics.const.inferred] Where a const argument is expected, an (optionally surrounded by any number of matching parentheses), called the *inferred const* ([path rules][paths.expr.complex-const-params], [array expression rules][expr.array.length-restriction]), can be used instead. This asks the compiler to infer the const argument if possible based on surrounding information.

r[items.generics.const.inferred.constraint] The inferred const cannot be used in item signatures.

r[items.generics.const.type-ambiguity] When there is ambiguity if a generic argument could be resolved as either a type or const argument, it is always resolved as a type. Placing the argument in a block expression can force it to be interpreted as a const argument.

```
type N = u32;
struct Foo<const N: usize>;
// The following is an error, because `N` is interpreted as the
type alias `N`.
fn foo<const N: usize>() -> Foo<N> { todo!() } // ERROR
// Can be fixed by wrapping in braces to force it to be
interpreted as the `N`
// const parameter:
fn bar<const N: usize>() -> Foo<{ N }> { todo!() } // ok
```

r[items.generics.const.variance] Unlike type and lifetime parameters, const parameters can be declared without being used inside of a parameterized item, with the exception of implementations as described in <u>generic implementations</u>:

```
// ok
struct Foo<const N: usize>;
enum Bar<const M: usize> { A, B }
```

```
// ERROR: unused parameter
struct Baz<T>;
struct Biz<'a>;
struct Unconstrained;
impl<const N: usize> Unconstrained {}
```

r[items.generics.const.exhaustiveness] When resolving a trait bound obligation, the exhaustiveness of all implementations of const parameters is not considered when determining if the bound is satisfied. For example, in the following, even though all possible const values for the bool type are implemented, it is still an error that the trait bound is not satisfied:

```
struct Foo<const B: bool>;
trait Bar {}
impl Bar for Foo<true> {}
impl Bar for Foo<false> {}
fn needs_bar(_: impl Bar) {}
fn generic<const B: bool>() {
    let v = Foo::<B>;
```

```
needs_bar(v); // ERROR: trait bound `Foo<B>: Bar` is not
satisfied
}
```

r[items.generics.where]

#### Where clauses

r[items.generics.where.syntax]
WhereClause -> `where` ( WhereClauseItem `,` )\*
WhereClauseItem?

WhereClauseItem ->

LifetimeWhereClauseItem

| TypeBoundWhereClauseItem

```
LifetimeWhereClauseItem -> Lifetime `:` LifetimeBounds
```

TypeBoundWhereClauseItem -> ForLifetimes? Type `:` TypeParamBounds?

r[items.generics.where.intro] *Where clauses* provide another way to specify bounds on type and lifetime parameters as well as a way to specify bounds on types that aren't type parameters.

r[items.generics.where.higher-ranked-lifetimes] The for keyword can be used to introduce <u>higher-ranked lifetimes</u>. It only allows [LifetimeParam] parameters.

```
struct A<T>
where
    T: Iterator, // Could use A<T: Iterator>
instead
    T::Item: Copy, // Bound on an associated type
    String: PartialEq<T>, // Bound on `String`, using the
type parameter
    i32: Default, // Allowed, but not useful
{
    f: T,
}
```

r[items.generics.attributes]

## Attributes

Generic lifetime and type parameters allow <u>attributes</u> on them. There are no built-in attributes that do anything in this position, although custom derive attributes may give meaning to it.

This example shows using a custom derive attribute to modify the meaning of a generic parameter.

```
// Assume that the derive for MyFlexibleClone declared
`my_flexible_clone` as
// an attribute it understands.
#[derive(MyFlexibleClone)]
struct Foo<#[my_flexible_clone(unbounded)] H> {
    a: *const H
}
```

r[items.associated]

# **Associated Items**

```
r[items.associated.syntax]
AssociatedItem ->
OuterAttribute* (
    MacroInvocationSemi
    | ( Visibility? ( TypeAlias | ConstantItem | Function ) )
)
```

r[items.associated.intro] *Associated Items* are the items declared in <u>traits</u> or defined in <u>implementations</u>. They are called this because they are defined on an associate type — the type in the implementation.

r[items.associated.kinds] They are a subset of the kinds of items you can declare in a module. Specifically, there are <u>associated functions</u> (including methods), <u>associated types</u>, and <u>associated constants</u>.

r[items.associated.related] Associated items are useful when the associated item is logically related to the associating item. For example, the is\_some method on Option is intrinsically related to Options, so should be associated.

r[items.associated.decl-def] Every associated item kind comes in two varieties: definitions that contain the actual implementation and declarations that declare signatures for definitions.

r[items.associated.trait-items] It is the declarations that make up the contract of traits and what is available on generic types.

r[items.associated.fn]

### **Associated functions and methods**

r[items.associated.fn.intro] *Associated functions* are <u>functions</u> associated with a type.

r[items.associated.fn.decl] An *associated function declaration* declares a signature for an associated function definition. It is written as a function item, except the function body is replaced with a ;.

r[items.associated.name] The identifier is the name of the function.

r[items.associated.same-signature] The generics, parameter list, return type, and where clause of the associated function must be the same as the associated function declarations's.

r[items.associated.fn.def] An *associated function definition* defines a function associated with another type. It is written the same as a <u>function</u> <u>item</u>.

[!NOTE] A common example is an associated function named new that returns a value of the type with which it is associated.

```
struct Struct {
    field: i32
}
impl Struct {
    fn new() -> Struct {
        Struct {
            field: 0i32
            }
        }
fn main () {
        let _struct = Struct::new();
}
```

r[items.associated.fn.qualified-self] When the associated function is declared on a trait, the function can also be called with a <u>path</u> that is a path

to the trait appended by the name of the trait. When this happens, it is substituted for <\_ as Trait>::function\_name.

```
trait Num {
    fn from_i32(n: i32) -> Self;
}
impl Num for f64 {
    fn from_i32(n: i32) -> f64 { n as f64 }
}
// These 4 are all equivalent in this case.
let _: f64 = Num::from_i32(42);
let _: f64 = <_ as Num>::from_i32(42);
let _: f64 = <f64 as Num>::from_i32(42);
let _: f64 = f64::from_i32(42);
```

r[items.associated.fn.method]

#### **Methods**

r[items.associated.fn.method.intro] Associated functions whose first parameter is named self are called *methods* and may be invoked using the <u>method call operator</u>, for example,  $\times.foo()$ , as well as the usual function call notation.

r[items.associated.fn.method.self-ty] If the type of the self parameter is specified, it is limited to types resolving to one generated by the following grammar (where 'lt denotes some arbitrary lifetime):

```
P = &'lt S | &'lt mut S | Box<S> | Rc<S> | Arc<S> | Pin<P>
S = Self | P
```

The Self terminal in this grammar denotes a type resolving to the implementing type. This can also include the contextual type alias Self, other type aliases, or associated type projections resolving to the implementing type.

```
# use std::rc::Rc;
# use std::sync::Arc;
```

```
# use std::pin::Pin;
// Examples of methods implemented on struct `Example`.
struct Example;
type Alias = Example;
trait Trait { type Output; }
impl Trait for Example { type Output = Example; }
impl Example {
    fn by_value(self: Self) {}
    fn by_ref(self: &Self) {}
    fn by_ref_mut(self: &mut Self) {}
    fn by_box(self: Box<Self>) {}
    fn by_rc(self: Rc<Self>) {}
    fn by_arc(self: Arc<Self>) {}
    fn by_pin(self: Pin<&Self>) {}
    fn explicit_type(self: Arc<Example>) {}
    fn with_lifetime<'a>(self: &'a Self) {}
    fn nested<'a>(self: &mut &'a Arc<Rc<Box<Alias>>>) {}
    fn via_projection(self: <Example as Trait>::Output) {}
}
```

r[associated.fn.method.self-pat-shorthands] Shorthand syntax can be used without specifying a type, which have the following equivalents:

Shorthand	Equivalent
self	self: Self
&'lifetime self	self: &'lifetime Self
&'lifetime mut self	self: &'lifetime mut Self

[!NOTE] Lifetimes can be, and usually are, elided with this shorthand.

r[associated.fn.method.self-pat-mut] If the self parameter is prefixed with mut, it becomes a mutable variable, similar to regular parameters using a mut <u>identifier pattern</u>. For example:

```
trait Changer: Sized {
   fn change(mut self) {}
   fn modify(mut self: Box<Self>) {}
}
```

As an example of methods on a trait, consider the following:

```
# type Surface = i32;
# type BoundingBox = i32;
trait Shape {
    fn draw(&self, surface: Surface);
    fn bounding_box(&self) -> BoundingBox;
}
```

This defines a trait with two methods. All values that have <u>implementations</u> of this trait while the trait is in scope can have their draw and bounding\_box methods called.

```
# type Surface = i32;
# type BoundingBox = i32;
# trait Shape {
      fn draw(&self, surface: Surface);
#
      fn bounding_box(&self) -> BoundingBox;
#
# }
#
struct Circle {
    // ...
}
impl Shape for Circle {
    // ...
    fn draw(&self, _: Surface) {}
#
    fn bounding_box(&self) -> BoundingBox { 0i32 }
#
}
# impl Circle {
      fn new() -> Circle { Circle{} }
#
# }
```

```
#
let circle_shape = Circle::new();
let bounding_box = circle_shape.bounding_box();
```

r[items.associated.fn.params.edition2018]

[!EDITION-2018] In the 2015 edition, it is possible to declare trait methods with anonymous parameters (e.g. fn foo(u8)). This is deprecated and an error as of the 2018 edition. All parameters must have an argument name.

r[items.associated.fn.param-attributes]

#### Attributes on method parameters

Attributes on method parameters follow the same rules and restrictions as <u>regular function parameters</u>.

r[items.associated.type]

# **Associated Types**

r[items.associated.type.intro] *Associated types* are <u>type aliases</u> associated with another type.

r[items.associated.type.restrictions] Associated types cannot be defined in <u>inherent implementations</u> nor can they be given a default implementation in traits.

r[items.associated.type.decl] An *associated type declaration* declares a signature for associated type definitions. It is written in one of the following forms, where Assoc is the name of the associated type, Params is a comma-separated list of type, lifetime or const parameters, Bounds is a plus-separated list of trait bounds that the associated type must meet, and WhereBounds is a comma-separated list of bounds that the parameters must meet:

type Assoc; type Assoc: Bounds; type Assoc<Params>; type Assoc<Params>: Bounds; type Assoc<Params> where WhereBounds; type Assoc<Params>: Bounds where WhereBounds;

r[items.associated.type.name] The identifier is the name of the declared type alias.

r[items.associated.type.impl-fulfillment] The optional trait bounds must be fulfilled by the implementations of the type alias.

r[items.associated.type.sized] There is an implicit <u>Sized</u> bound on associated types that can be relaxed using the special <u>?Sized</u> bound.

r[items.associated.type.def] An *associated type definition* defines a type alias for the implementation of a trait on a type

r[items.associated.type.def.restriction] They are written similarly to an *associated type declaration*, but cannot contain Bounds, but instead must contain a Type:

```
type Assoc = Type;
type Assoc<Params> = Type; // the type `Type` here may
```

```
reference `Params`
type Assoc<Params> = Type where WhereBounds;
type Assoc<Params> where WhereBounds = Type; // deprecated,
prefer the form above
```

r[items.associated.type.alias] If a type Item has an associated type Assoc from a trait Trait, then <Item as Trait>::Assoc is a type that is an alias of the type specified in the associated type definition

r[items.associated.type.param] Furthermore, if Item is a type parameter, then Item::Assoc can be used in type parameters.

r[items.associated.type.generic] Associated types may include <u>generic</u> <u>parameters</u> and <u>where clauses</u>; these are often referred to as <u>generic</u> associated types, or GATs. If the type Thing has an associated type Item from a trait Trait with the generics <'a> , the type can be named like <Thing as Trait>::Item<'x>, where 'x is some lifetime in scope. In this case, 'x will be used wherever 'a appears in the associated type definitions on impls.

```
trait AssociatedType {
    // Associated type declaration
    type Assoc;
}
struct Struct;
impl AssociatedType for Struct {
    // AssociatedType for Struct {
    // Associated type definition
    type Assoc = OtherStruct;
}
impl OtherStruct {
    fn new() -> OtherStruct {
        OtherStruct
    }
```

fn main() {
 // Usage of the associated type to refer to OtherStruct as
<Struct as AssociatedType>::Assoc
 let \_other\_struct: OtherStruct = <Struct as
AssociatedType>::Assoc::new();
}

An example of associated types with generics and where clauses: struct ArrayLender<'a, T>(&'a mut [T; 16]);

}

```
trait Lend {
   // Generic associated type declaration
   type Lender<'a> where Self: 'a;
   fn lend<'a>(&'a mut self) -> Self::Lender<'a>;
}
impl<T> Lend for [T; 16] {
   // Generic associated type definition
   type Lender<'a> = ArrayLender<'a, T> where Self: 'a;
   fn lend<'a>(&'a mut self) -> Self::Lender<'a> {
       ArrayLender(self)
   }
}
    borrow<'a, T: Lend>(array: &'a mut T) -> <T
fn
                                                            as
Lend>::Lender<'a> {
   array.lend()
}
fn main() {
   let mut array = [Ousize; 16];
   let lender = borrow(&mut array);
}
```

# **Associated Types Container Example**

Consider the following example of a **Container** trait. Notice that the type is available for use in the method signatures:

```
trait Container {
   type E;
   fn empty() -> Self;
   fn insert(&mut self, elem: Self::E);
}
```

In order for a type to implement this trait, it must not only provide implementations for every method, but it must specify the type E. Here's an implementation of Container for the standard library type Vec:

```
# trait Container {
# type E;
# fn empty() -> Self;
# fn insert(&mut self, elem: Self::E);
# }
impl<T> Container for Vec<T> {
   type E = T;
   fn empty() -> Vec<T> { Vec::new() }
   fn insert(&mut self, x: T) { self.push(x); }
}
```

#### **Relationship between Bounds and WhereBounds**

```
In this example:
# use std::fmt::Debug;
trait Example {
    type Output<T>: Ord where T: Debug;
}
```

Given a reference to the associated type like <X as Example>::Output<Y>, the associated type itself must be Ord, and the type Y must be Debug.

```
r[items.associated.type.generic-where-clause]
```

# Required where clauses on generic associated types

r[items.associated.type.generic-where-clause.intro] Generic associated type declarations on traits currently may require a list of where clauses, dependent on functions in the trait and how the GAT is used. These rules may be loosened in the future; updates can be found <u>on the generic associated types initiative repository</u>.

r[items.associated.type.generic-where-clause.valid-fn] In a few words, these where clauses are required in order to maximize the allowed definitions of the associated type in impls. To do this, any clauses that *can be proven to hold* on functions (using the parameters of the function or trait) where a GAT appears as an input or output must also be written on the GAT itself.

```
trait LendingIterator {
   type Item<'x> where Self: 'x;
   fn next<'a>(&'a mut self) -> Self::Item<'a>;
}
```

In the above, on the next function, we can prove that Self: 'a, because of the implied bounds from &'a mut self; therefore, we must write the equivalent bound on the GAT itself: where Self: 'x.

r[items.associated.type.generic-where-clause.intersection] When there are multiple functions in a trait that use the GAT, then the *intersection* of the bounds from the different functions are used, rather than the union.

```
trait Check<T> {
   type Checker<'x>;
   fn create_checker<'a>(item: &'a T) -> Self::Checker<'a>;
   fn do_check(checker: Self::Checker<'_>);
}
```

In this example, no bounds are required on the type Checker<'a>;. While we know that T: 'a on create\_checker, we do not know that on do\_check. However, if do\_check was commented out, then the where T: 'x bound would be required on Checker. r[items.associated.type.generic-where-clause.forward] The bounds on associated types also propagate required where clauses.

```
trait Iterable {
   type Item<'a> where Self: 'a;
   type Iterator<'a>: Iterator<Item = Self::Item<'a>> where
Self: 'a;
   fn iter<'a>(&'a self) -> Self::Iterator<'a>;
}
```

Here, where Self: 'a is required on Item because of iter. However, Item is used in the bounds of Iterator, the where Self: 'a clause is also required there.

r[items.associated.type.generic-where-clause.static] Finally, any explicit uses of 'static on GATs in the trait do not count towards the required bounds.

```
trait StaticReturn {
   type Y<'a>;
   fn foo(&self) -> Self::Y<'static>;
}
```

r[items.associated.const]

#### **Associated Constants**

r[items.associated.const.intro] *Associated constants* are <u>constants</u> associated with a type.

r[items.associated.const.decl] An *associated constant declaration* declares a signature for associated constant definitions. It is written as const, then an identifier, then : , then a type, finished by a ; .

r[items.associated.const.name] The identifier is the name of the constant used in the path. The type is the type that the definition has to implement.

r[items.associated.const.def] An *associated constant definition* defines a constant associated with a type. It is written the same as a <u>constant item</u>.

r[items.associated.const.eval] Associated constant definitions undergo <u>constant evaluation</u> only when referenced. Further, definitions that include <u>generic parameters</u> are evaluated after monomorphization.

```
struct Struct;
struct GenericStruct<const ID: i32>;
impl Struct {
    // Definition not immediately evaluated
   const PANIC: () = panic!("compile-time panic");
}
impl<const ID: i32> GenericStruct<ID> {
    // Definition not immediately evaluated
    const NON_ZERO: () = if ID == 0 {
        panic!("contradiction")
    };
}
fn main() {
    // Referencing Struct::PANIC causes compilation error
    let _ = Struct::PANIC;
    // Fine, ID is not 0
```
```
let _ = GenericStruct::<1>::NON_ZERO;
// Compilation error from evaluating NON_ZERO with ID=0
let _ = GenericStruct::<0>::NON_ZERO;
}
```

#### **Associated Constants Examples**

```
A basic example:
trait ConstantId {
    const ID: i32;
}
struct Struct;
impl ConstantId for Struct {
    const ID: i32 = 1;
}
fn main() {
    assert_eq!(1, Struct::ID);
}
 Using default values:
trait ConstantIdDefault {
    const ID: i32 = 1;
}
struct Struct;
struct OtherStruct;
impl ConstantIdDefault for Struct {}
impl ConstantIdDefault for OtherStruct {
    const ID: i32 = 5;
}
```

```
fn main() {
    assert_eq!(1, Struct::ID);
    assert_eq!(5, OtherStruct::ID);
}
```

r[attributes]

## Attributes

r[attributes.syntax]

InnerAttribute -> `#` `!` `[` Attr `]`

OuterAttribute -> `#` `[` Attr `]`

Attr ->

SimplePath AttrInput?

| `unsafe` `(` SimplePath AttrInput? `)`

AttrInput ->

DelimTokenTree

| `=` Expression

r[attributes.intro] An *attribute* is a general, free-form metadatum that is interpreted according to name, convention, language, and compiler version. Attributes are modeled on Attributes in <u>ECMA-335</u>, with the syntax coming from <u>ECMA-334</u> (C#).

r[attributes.inner] *Inner attributes*, written with a bang (!) after the hash (#), apply to the item that the attribute is declared within. *Outer attributes*, written without the bang after the hash, apply to the thing that follows the attribute.

r[attributes.input] The attribute consists of a path to the attribute, followed by an optional delimited token tree whose interpretation is defined by the attribute. Attributes other than macro attributes also allow the input to be an equals sign (=) followed by an expression. See the <u>meta item</u> <u>syntax</u> below for more details.

r[attributes.safety] An attribute may be unsafe to apply. To avoid undefined behavior when using these attributes, certain obligations that cannot be checked by the compiler must be met. To assert these have been, the attribute is wrapped in unsafe(..), e.g. #[unsafe(no\_mangle)].

The following attributes are unsafe:

- <u>export name</u>
- <u>link section</u>
- <u>naked</u>
- <u>no mangle</u>

r[attributes.kind] Attributes can be classified into the following kinds:

- Built-in attributes
- <u>Proc macro attributes</u>
- <u>Derive macro helper attributes</u>
- <u>Tool attributes</u>

r[attributes.allowed-position] Attributes may be applied to many things in the language:

- All <u>item declarations</u> accept outer attributes while <u>external blocks</u>, <u>functions</u>, <u>implementations</u>, and <u>modules</u> accept inner attributes.
- Most <u>statements</u> accept outer attributes (see <u>Expression Attributes</u> for limitations on expression statements).
- <u>Block expressions</u> accept outer and inner attributes, but only when they are the outer expression of an <u>expression statement</u> or the final expression of another block expression.
- <u>Enum</u> variants and <u>struct</u> and <u>union</u> fields accept outer attributes.
- <u>Match expression arms</u> accept outer attributes.
- <u>Generic lifetime or type parameter</u> accept outer attributes.
- Expressions accept outer attributes in limited situations, see <u>Expression Attributes</u> for details.
- <u>Function</u>, <u>closure</u> and <u>function pointer</u> parameters accept outer attributes. This includes attributes on variadic parameters denoted with ... in function pointers and <u>external blocks</u>.

Some examples of attributes:

```
// General metadata applied to the enclosing module or crate.
#![crate_type = "lib"]
```

```
// A function marked as a unit test
#[test]
```

```
fn test_foo() {
   /* ... */
}
// A conditionally-compiled module
#[cfg(target_os = "linux")]
mod bar {
   /* ... */
}
// A lint attribute used to suppress a warning/error
#[allow(non_camel_case_types)]
type int8_t = i8;
// Inner attribute applies to the entire function.
fn some_unused_variables() {
  #![allow(unused_variables)]
  let x = ();
  let y = ();
  let z = ();
}
 r[attributes.meta]
```

#### **Meta Item Attribute Syntax**

r[attributes.meta.intro] A "meta item" is the syntax used for the [Attr] rule by most <u>built-in attributes</u>. It has the following grammar:

r[attributes.meta.literal-expr] Expressions in meta items must macroexpand to literal expressions, which must not include integer or float type suffixes. Expressions which are not literal expressions will be syntactically accepted (and can be passed to proc-macros), but will be rejected after parsing.

r[attributes.meta.order] Note that if the attribute appears within another macro, it will be expanded after that outer macro. For example, the following code will expand the Serialize proc-macro first, which must preserve the include\_str! call in order for it to be expanded:

```
#[derive(Serialize)]
struct Foo {
    #[doc = include_str!("x.md")]
    x: u32
}
```

r[attributes.meta.order-macro] Additionally, macros in attributes will be expanded only after all other attributes applied to the item:

```
#[macro_attr1] // expanded first
#[doc = mac!()] // `mac!` is expanded fourth.
```

```
#[macro_attr2] // expanded second
#[derive(MacroDerive1, MacroDerive2)] // expanded third
fn foo() {}
```

r[attributes.meta.builtin] Various built-in attributes use different subsets of the meta item syntax to specify their inputs. The following grammar rules show some commonly used forms:

```
r[attributes.meta.builtin.syntax]
@root MetaWord ->
IDENTIFIER
MetaNameValueStr ->
IDENTIFIER `=` (STRING_LITERAL | RAW_STRING_LITERAL)
@root MetaListPaths ->
IDENTIFIER `(` ( SimplePath (`,` SimplePath)* `,`? )? `)`
@root MetaListIdents ->
IDENTIFIER `(` ( IDENTIFIER (`,` IDENTIFIER)* `,`? )? `)`
@root MetaListNameValueStr ->
```

```
IDENTIFIER `(` ( MetaNameValueStr (`,` MetaNameValueStr)*
`,`? )? `)`
```

Some examples of meta items are:

Style	Example
[MetaWord]	no_std
[MetaNameValueStr]	<pre>doc = "example"</pre>
[MetaListPaths]	<pre>allow(unused, clippy::inline_always)</pre>
[MetaListIdents]	<pre>macro_use(foo, bar)</pre>
[MetaListNameValueStr]	<pre>link(name = "CoreFoundation", kind = "framework")</pre>

r[attributes.activity]

#### **Active and inert attributes**

r[attributes.activity.intro] An attribute is either active or inert. During attribute processing, *active attributes* remove themselves from the thing they are on while *inert attributes* stay on.

The <u>cfg</u> and <u>cfg attr</u> attributes are active. <u>Attribute macros</u> are active. All other attributes are inert.

r[attributes.tool]

#### **Tool attributes**

r[attributes.tool.intro] The compiler may allow attributes for external tools where each tool resides in its own module in the <u>tool prelude</u>. The first segment of the attribute path is the name of the tool, with one or more additional segments whose interpretation is up to the tool.

r[attributes.tool.ignored] When a tool is not in use, the tool's attributes are accepted without a warning. When the tool is in use, the tool is responsible for processing and interpretation of its attributes.

r[attributes.tool.prelude] Tool attributes are not available if the <u>no implicit prelude</u> attribute is used.

```
// Tells the rustfmt tool to not format the following element.
#[rustfmt::skip]
struct S {
}
// Controls the "cyclomatic complexity" threshold for the
clippy tool.
#[clippy::cyclomatic_complexity = "100"]
pub fn f() {}
```

```
[!NOTE] rustc currently recognizes the tools "clippy", "rustfmt", "diagnostic", "miri" and "rust_analyzer".
```

r[attributes.builtin]

### **Built-in attributes index**

The following is an index of all built-in attributes.

- Conditional compilation
  - <u>cfg</u> --- Controls conditional compilation.
  - <u>cfg attr</u> --- Conditionally includes attributes.
- Testing
  - <u>test</u> --- Marks a function as a test.
  - <u>ignore</u> --- Disables a test function.
  - <u>should panic</u> --- Indicates a test should generate a panic.
- Derive
  - derive --- Automatic trait implementations.
  - <u>automatically derived</u> --- Marker for implementations created by derive.
- Macros
  - <u>macro export</u> --- Exports a <u>macro\_rules</u> macro for cross-crate usage.
  - <u>macro use</u> --- Expands macro visibility, or imports macros from other crates.
  - proc macro --- Defines a function-like macro.
  - proc macro derive --- Defines a derive macro.
  - proc macro attribute --- Defines an attribute macro.
- Diagnostics
  - <u>allow</u>, <u>expect</u>, <u>warn</u>, <u>deny</u>, <u>forbid</u> --- Alters the default lint level.
  - <u>deprecated</u> --- Generates deprecation notices.
  - <u>must use</u> --- Generates a lint for unused values.

- <u>diagnostic::on unimplemented</u> --- Hints the compiler to emit a certain error message if a trait is not implemented.
- <u>diagnostic::do not recommend</u> --- Hints the compiler to not show a certain trait impl in error messages.
- ABI, linking, symbols, and FFI
  - <u>link</u> --- Specifies a native library to link with an extern block.
  - <u>link name</u> --- Specifies the name of the symbol for functions or statics in an extern block.
  - <u>link ordinal</u> --- Specifies the ordinal of the symbol for functions or statics in an extern block.
  - <u>no link</u> --- Prevents linking an extern crate.
  - <u>repr</u> --- Controls type layout.
  - <u>crate type</u> --- Specifies the type of crate (library, executable, etc.).
  - <u>no main</u> --- Disables emitting the main symbol.
  - <u>export name</u> --- Specifies the exported symbol name for a function or static.
  - <u>link section</u> --- Specifies the section of an object file to use for a function or static.
  - <u>no mangle</u> --- Disables symbol name encoding.
  - <u>used</u> --- Forces the compiler to keep a static item in the output object file.
  - <u>crate name</u> --- Specifies the crate name.
- Code generation
  - <u>inline</u> --- Hint to inline code.
  - **cold** --- Hint that a function is unlikely to be called.
  - <u>naked</u> --- Prevent the compiler from emitting a function prologue.
  - <u>no builtins</u> --- Disables use of certain built-in functions.
  - <u>target feature</u> --- Configure platform-specific code generation.

- track caller --- Pass the parent call location to
  std::panic::Location::caller().
- <u>instruction set</u> --- Specify the instruction set used to generate a functions code
- Documentation
  - doc --- Specifies documentation. See <u>The Rustdoc Book</u> for more information. <u>Doc comments</u> are transformed into doc attributes.
- Preludes
  - <u>no std</u> --- Removes std from the prelude.
  - <u>no implicit prelude</u> --- Disables prelude lookups within a module.
- Modules
  - <u>path</u> --- Specifies the filename for a module.
- Limits
  - <u>recursion limit</u> --- Sets the maximum recursion limit for certain compile-time operations.
  - <u>type length limit</u> --- Sets the maximum size of a polymorphic type.
- Runtime
  - <u>panic handler</u> --- Sets the function to handle panics.
  - <u>global allocator</u> --- Sets the global memory allocator.
  - <u>windows subsystem</u> --- Specifies the windows subsystem to link with.
- Features
  - feature --- Used to enable unstable or experimental compiler features. See <u>The Unstable Book</u> for features implemented in rustc.

- Type System
  - <u>non exhaustive</u> --- Indicate that a type will have more fields/variants added in future.
- Debugger
  - <u>debugger visualizer</u> --- Embeds a file that specifies debugger output for a type.
  - <u>collapse debuginfo</u> --- Controls how macro invocations are encoded in debuginfo.

r[attributes.testing]

# **Testing attributes**

The following <u>attributes</u> are used for specifying functions for performing tests. Compiling a crate in "test" mode enables building the test functions along with a test harness for executing the tests. Enabling the test mode also enables the <u>test conditional compilation option</u>.

r[attributes.testing.test]

#### The test attribute

r[attributes.testing.test.intro] The *test* <u>attribute</u> marks a function to be executed as a test.

```
[!EXAMPLE]
```

```
# pub fn add(left: u64, right: u64) -> u64 { left + right }
#[test]
fn it_works() {
    let result = add(2, 2);
    assert_eq!(result, 4);
}
```

r[attributes.testing.test.syntax] The test attribute uses the [MetaWord] syntax and thus does not take any inputs.

r[attributes.testing.test.allowed-positions] The test attribute may only be applied to <u>free functions</u> that are monomorphic, that take no arguments, and where the return type implements the <u>Termination</u> trait.

[!NOTE] Some of types that implement the <u>Termination</u> trait include:

```
• ()
```

```
• Result<T, E> where T: Termination, E: Debug
```

r[attributes.testing.test.duplicates] Only the first instance of test on a function is honored.

[!NOTE] Subsequent test attributes are currently ignored and rustc warns about these.

r[attributes.testing.test.stdlib] The test attribute is exported from the standard library prelude as [std::prelude::v1::test].

r[attributes.testing.test.enabled] These functions are only compiled when in test mode.

[!NOTE] The test mode is enabled by passing the --test argument to rustc or using cargo test.

r[attributes.testing.test.success] The test harness calls the returned value's <u>report</u> method, and classifies the test as passed or failed depending on whether the resulting <u>ExitCode</u> represents successful termination. In particular:

- Tests that return () pass as long as they terminate and do not panic.
- Tests that return a Result<(), E> pass as long as they return Ok(()).
- Tests that return ExitCode::SUCCESS pass, and tests that return ExitCode::FAILURE fail.
- Tests that do not terminate neither pass nor fail.

```
[!EXAMPLE]
```

```
# use std::io;
# fn setup_the_thing() -> io::Result<i32> { 0k(1) }
# fn do_the_thing(s: &i32) -> io::Result<()> { 0k(()) }
#[test]
fn test_the_thing() -> io::Result<()> {
    let state = setup_the_thing()?; // expected to succeed
    do_the_thing(&state)?; // expected to succeed
    ok(())
}
```

r[attributes.testing.ignore]

#### The ignore attribute

r[attributes.testing.ignore.intro] The *ignore* <u>attribute</u> can be used with the [test attribute][attributes.testing.test] to tell the test harness to not execute that function as a test.

```
[!EXAMPLE]
#[test]
#[ignore]
fn check_thing() {
    // ...
}
```

[!NOTE] The rustc test harness supports the --include-ignored flag to force ignored tests to be run.

r[attributes.testing.ignore.syntax] The ignore attribute uses either the [MetaWord] or [MetaNameValueStr] syntax.

r[attributes.testing.ignore.reason] The [MetaNameValueStr] form of the ignore attribute provides a way to specify a reason why the test is ignored.

```
[!EXAMPLE]
#[test]
#[ignore = "not yet implemented"]
fn mytest() {
    // ...
}
```

r[attributes.testing.ignore.allowed-positions] The ignore attribute may be applied to functions annotated with the test attribute.

[!NOTE] rustc currently warns when ignore is used in some other situations. This may become an error in the future.

r[attributes.testing.ignore.duplicates] Only the first instance of ignore on a function is honored.

[!NOTE] rustc currently ignores duplicate ignore attributes. This may become an error in the future.

r[attributes.testing.ignore.behavior] Ignored tests are still compiled when in test mode, but they are not executed.

r[attributes.testing.should\_panic]

#### The should\_panic attribute

r[attributes.testing.should\_panic.intro] The *should\_panic* <u>attribute</u> causes a test to pass only if the [test function][attributes.testing.test] to which the attribute is applied panics.

```
[!EXAMPLE]
#[test]
#[should_panic(expected = "values don't match")]
fn mytest() {
    assert_eq!(1, 2, "values don't match");
}
```

r[attributes.testing.should\_panic.syntax] The should\_panic attribute
has one of the following forms:

```
• [MetaWord]
```

```
[!EXAMPLE]
#[test]
#[should_panic]
fn mytest() { panic!("error: some message, and more");
}
```

• [MetaNameValueStr] --- The given string must appear within the panic message for the test to pass.

```
[!EXAMPLE]
#[test]
#[should_panic = "some message"]
fn mytest() { panic!("error: some message, and more");
}
```

• [MetaListNameValueStr] --- As with the [MetaNameValueStr] syntax, the given string must appear within the panic message.

```
[!EXAMPLE]
#[test]
#[should_panic(expected = "some message")]
```

```
fn mytest() { panic!("error: some message, and more");
}
```

[!NOTE] rustc currently accepts the [MetaListNameValueStr] form with invalid syntax between the parentheses and emits a future-compatibility warning. This may

become a hard error in the future.

r[attributes.testing.should\_panic.allowed-positions] The should\_panic attribute may only be applied to functions annotated with the test attribute.

[!NOTE] **rustc** currently accepts this attribute in other positions with a warning. This may become a hard error in the future.

r[attributes.testing.should\_panic.duplicates] Only the first instance of should\_panic on a function is honored.

[!NOTE] **rustc** currently ignores subsequent **should\_panic** attributes and emits a future-compatibility warning. This may become a hard error in the future.

r[attributes.testing.should\_panic.expected] When the [MetaNameValueStr] form or the [MetaListNameValueStr] form with the expected key is used, the given string must appear somewhere within the panic message for the test to pass.

r[attributes.testing.should\_panic.return] The return type of the test function must be ().

r[attributes.derive]

## Derive

r[attributes.derive.intro] The *derive attribute* allows new <u>items</u> to be automatically generated for data structures.

r[attributes.derive.syntax] It uses the [MetaListPaths] syntax to specify a list of traits to implement or paths to <u>derive macros</u> to process.

For example, the following will create an <u>impl\_item</u> for the [PartialEq] and [Clone] traits for Foo, and the type parameter T will be given the PartialEq or Clone constraints for the appropriate impl:

```
#[derive(PartialEq, Clone)]
struct Foo<T> {
    a: i32,
    b: T,
}
```

The generated impl for PartialEq is equivalent to

```
# struct Foo<T> { a: i32, b: T }
impl<T: PartialEq> PartialEq for Foo<T> {
    fn eq(&self, other: &Foo<T>) -> bool {
        self.a == other.a && self.b == other.b
    }
}
```

r[attributes.derive.proc-macro] You can implement derive for your own traits through procedural macros.

r[attributes.derive.automatically\_derived]

#### The automatically\_derived attribute

The *automatically\_derived attribute* is automatically added to <u>implementations</u> created by the derive attribute for built-in traits. It has no direct effect, but it may be used by tools and diagnostic lints to detect these automatically generated implementations.

r[attributes.diagnostics]

# **Diagnostic attributes**

The following <u>attributes</u> are used for controlling or generating diagnostic messages during compilation.

r[attributes.diagnostics.lint]

#### Lint check attributes

A lint check names a potentially undesirable coding pattern, such as unreachable code or omitted documentation.

r[attributes.diagnostics.lint.level] The lint attributes allow, expect, warn, deny, and forbid use the [MetaListPaths] syntax to specify a list of lint names to change the lint level for the entity to which the attribute applies.

For any lint check C:

r[attributes.diagnostics.lint.allow]

 #[allow(C)] overrides the check for C so that violations will go unreported.

r[attributes.diagnostics.lint.expect]

#[expect(C)] indicates that lint C is expected to be emitted. The attribute will suppress the emission of C or issue a warning, if the expectation is unfulfilled.

r[attributes.diagnostics.lint.warn]

- #[warn(C)] warns about violations of C but continues compilation.
   r[attributes.diagnostics.lint.deny]
- #[deny(C)] signals an error after encountering a violation of C,
   r[attributes.diagnostics.lint.forbid]
- #[forbid(C)] is the same as deny(C), but also forbids changing the lint level afterwards,

[!NOTE] The lint checks supported by **rustc** can be found via **rustc** -W help, along with their default settings and are documented in the <u>rustc book</u>.

```
pub mod m1 {
```

// Missing documentation is ignored here

```
#[allow(missing_docs)]
pub fn undocumented_one() -> i32 { 1 }
// Missing documentation signals a warning here
#[warn(missing_docs)]
pub fn undocumented_too() -> i32 { 2 }
// Missing documentation signals an error here
#[deny(missing_docs)]
pub fn undocumented_end() -> i32 { 3 }
```

}

r[attributes.diagnostics.lint.override] Lint attributes can override the level specified from a previous attribute, as long as the level does not attempt to change a forbidden lint (except for deny, which is allowed inside a forbid context, but ignored). Previous attributes are those from a higher level in the syntax tree, or from a previous attribute on the same entity as listed in left-to-right source order.

This example shows how one can use allow and warn to toggle a particular check on and off:

```
#[warn(missing_docs)]
pub mod m2 {
    #[allow(missing_docs)]
    pub mod nested {
        // Missing documentation is ignored here
        pub fn undocumented_one() -> i32 { 1 }
        // Missing documentation signals a warning here,
        // despite the allow above.
        #[warn(missing_docs)]
        pub fn undocumented_two() -> i32 { 2 }
    }
    // Missing documentation signals a warning here
```

```
pub fn undocumented_too() -> i32 { 3 }
```

}

This example shows how one can use forbid to disallow uses of allow or expect for that lint check:

```
#[forbid(missing_docs)]
pub mod m3 {
    // Attempting to toggle warning signals an error here
    #[allow(missing_docs)]
    /// Returns 2.
    pub fn undocumented_too() -> i32 { 2 }
```

}

[!NOTE] **rustc** allows setting lint levels on the <u>command-line</u>, and also supports <u>setting caps</u> on the lints that are reported.

r[attributes.diagnostics.lint.reason]

#### **Lint Reasons**

All lint attributes support an additional reason parameter, to give context why a certain attribute was added. This reason will be displayed as part of the lint message if the lint is emitted at the defined level.

```
// `keyword_idents` is allowed by default. Here we deny it to
// avoid migration of identifiers when we update the edition.
#![deny(
```

keyword\_idents,

```
reason = "we want to avoid these idents to be future
compatible"
```

)]

```
// This name was allowed in Rust's 2015 edition. We still aim
to avoid
```

```
// this to be future compatible and not confuse end users.
```

```
fn dyn() {}
```

Here is another example, where the lint is allowed with a reason:

```
use std::path::PathBuf;
pub fn get_path() -> PathBuf {
    // The `reason` parameter on `allow` attributes acts as
documentation for the reader.
    #[allow(unused_mut, reason = "this is only modified on
some platforms")]
    let mut file_name = PathBuf::from("git");
    #[cfg(target_os = "windows")]
    file_name.set_extension("exe");
    file_name
}
```

r[attributes.diagnostics.expect]

#### The #[expect] attribute

r[attributes.diagnostics.expect.intro] The #[expect(C)] attribute creates
a lint expectation for lint C. The expectation will be fulfilled, if a #
[warn(C)] attribute at the same location would result in a lint emission. If
the expectation is unfulfilled, because lint C would not be emitted, the
unfulfilled\_lint\_expectations lint will be emitted at the attribute.

```
fn main() {
```

// This `#[expect]` attribute creates a lint expectation,
that the `unused\_variables`

 $\ensuremath{//}$  lint would be emitted by the following statement. This expectation is

// unfulfilled, since the `question` variable is used by
the `println!` macro.

// Therefore, the `unfulfilled\_lint\_expectations` lint
will be emitted at the

// attribute.

#[expect(unused\_variables)]

let question = "who lives in a pineapple under the sea?";

println!("{question}");

}

r[attributes.diagnostics.expect.fulfillment] The lint expectation is only fulfilled by lint emissions which have been suppressed by the expect attribute. If the lint level is modified in the scope with other level attributes like allow or warn, the lint emission will be handled accordingly and the expectation will remain unfulfilled.

```
#[expect(unused_variables)]
```

```
fn select_song() {
```

// This will emit the `unused\_variables` lint at the warn level  $% \left( \left( {{{\left( {{\left( {{{\left( {{{}}} \right)}} \right)}} \right)}} \right)$ 

// as defined by the `warn` attribute. This will not
fulfill the

// expectation above the function.

#[warn(unused\_variables)]

let song\_name = "Crab Rave";

// The `allow` attribute suppresses the lint emission.
This will not

// fulfill the expectation as it has been suppressed by
the `allow`

// attribute and not the `expect` attribute above the
function.

#[allow(unused\_variables)]

let song\_creator = "Noisestorm";

```
// This `expect` attribute will suppress the
`unused_variables` lint emission
    // at the variable. The `expect` attribute above the
function will still not
    // be fulfilled, since this lint emission has been
suppressed by the local
    // expect attribute.
    #[expect(unused_variables)]
    let song_version = "Monstercat Release";
}
```

r[attributes.diagnostics.expect.independent] If the expect attribute contains several lints, each one is expected separately. For a lint group it's enough if one lint inside the group has been emitted:

```
// This expectation will be fulfilled by the unused value
inside the function
// since the emitted `unused_variables` lint is inside the
`unused` lint group.
#[expect(unused)]
pub fn thoughts() {
    let unused = "I'm running out of examples";
}
pub fn another_example() {
      // This attribute creates two lint expectations. The
`unused mut` lint will be
    // suppressed and with that fulfill the first expectation.
The `unused variables`
    // wouldn't be emitted, since the variable is used. That
expectation will therefore
    // be unsatisfied, and a warning will be emitted.
    #[expect(unused mut, unused variables)]
    let mut link = "https://www.rust-lang.org/";
    println!("Welcome to our community: {link}");
}
```

[!NOTE]Thebehaviorof#[expect(unfulfilled\_lint\_expectations)]is currently defined toalways generate theunfulfilled\_lint\_expectationslint.

r[attributes.diagnostics.lint.group]

### Lint groups

Lints may be organized into named groups so that the level of related lints can be adjusted together. Using a named group is equivalent to listing out the lints within that group.

```
// This allows all lints in the "unused" group.
#[allow(unused)]
// This overrides the "unused_must_use" lint from the "unused"
// group to deny.
#[deny(unused_must_use)]
fn example() {
        // This does
                       not
                             generate a warning
                                                   because
                                                            the
"unused variables"
    // lint is in the "unused" group.
    let x = 1;
    // This generates an error because the result is unused and
    // "unused_must_use" is marked as "deny".
         std::fs::remove_file("some_file"); // ERROR:
                                                         unused
`Result` that must be used
}
```

r[attributes.diagnostics.lint.group.warnings] There is a special group named "warnings" which includes all lints at the "warn" level. The "warnings" group ignores attribute order and applies to all lints that would otherwise warn within the entity.

```
# unsafe fn an_unsafe_fn() {}
// The order of these two attributes does not matter.
#[deny(warnings)]
// The unsafe_code lint is normally "allow" by default.
#[warn(unsafe_code)]
fn example_err() {
```

```
// This is an error because the `unsafe_code` warning has
// been lifted to "deny".
unsafe { an_unsafe_fn() } // ERROR: usage of `unsafe` block
}
```

r[attributes.diagnostics.lint.tool]

#### **Tool lint attributes**

r[attributes.diagnostics.lint.tool.intro] Tool lints allows using scoped lints, to allow, warn, deny or forbid lints of certain tools.

r[attributes.diagnostics.lint.tool.activation] Tool lints only get checked when the associated tool is active. If a lint attribute, such as allow, references a nonexistent tool lint, the compiler will not warn about the nonexistent lint until you use the tool.

Otherwise, they work just like regular lint attributes:

```
// set the entire `pedantic` clippy lint group to warn
#![warn(clippy::pedantic)]
// silence warnings from the `filter_map` clippy lint
#![allow(clippy::filter_map)]
fn main() {
    // ...
}
// silence the `cmp_nan` clippy lint just for this function
#[allow(clippy::cmp_nan)]
fn foo() {
    // ...
}
```

[!NOTE] **rustc** currently recognizes the tool lints for "<u>clippy</u>" and "<u>rustdoc</u>".

r[attributes.diagnostics.deprecated]
### The deprecated attribute

r[attributes.diagnostics.deprecated.intro] The deprecated attribute
marks an item as deprecated. rustc will issue warnings on usage of #
[deprecated] items. rustdoc will show item deprecation, including the
since version and note, if available.

r[attributes.diagnostics.deprecated.syntax] The deprecated attribute has several forms:

- deprecated --- Issues a generic message.
- deprecated = "message" --- Includes the given string in the deprecation message.
- [MetaListNameValueStr] syntax with two optional fields:
  - since --- Specifies a version number when the item was deprecated. rustc does not currently interpret the string, but external tools like <u>Clippy</u> may check the validity of the value.
  - note --- Specifies a string that should be included in the deprecation message. This is typically used to provide an explanation about the deprecation and preferred alternatives.

r[attributes.diagnostic.deprecated.allowed-positions] The deprecated attribute may be applied to any item, trait item, enum variant, struct field, external block item, or macro definition. It cannot be applied to trait implementation items. When applied to an item containing other items, such as a module or implementation, all child items inherit the deprecation attribute.

Here is an example:

```
#[deprecated(since = "5.2.0", note = "foo was rarely used.
Users should instead use bar")]
pub fn foo() {}
pub fn bar() {}
```

The <u>RFC</u> contains motivations and more details.

r[attributes.diagnostics.must\_use]

#### The must\_use attribute

r[attributes.diagnostics.must\_use.intro] The *must\_use* attribute is used to issue a diagnostic warning when a value is not "used".

r[attributes.diagnostics.must\_use.allowed-positions] The must\_use attribute can be applied to user-defined composite types (<u>structs</u>, <u>enums</u>, and <u>unions</u>), <u>functions</u>, and <u>traits</u>.

r[attributes.diagnostics.must\_use.message] The must\_use attribute may include a message by using the [MetaNameValueStr] syntax such as # [must\_use = "example message"]. The message will be given alongside the warning.

r[attributes.diagnostics.must\_use.type] When used on user-defined composite types, if the <u>expression</u> of an <u>expression statement</u> has that type, then the <u>unused\_must\_use</u> lint is violated.

```
#[must_use]
struct MustUse {
    // some fields
}
# impl MustUse {
# fn new() -> MustUse { MustUse {} }
# }
#
// Violates the `unused_must_use` lint.
MustUse::new();
```

r[attributes.diagnostics.must\_use.fn] When used on a function, if the <u>expression</u> of an <u>expression statement</u> is a <u>call expression</u> to that function, then the <u>unused\_must\_use</u> lint is violated.

```
#[must_use]
fn five() -> i32 { 5i32 }
// Violates the unused_must_use lint.
five();
```

r[attributes.diagnostics.must\_use.trait] When used on a <u>trait declaration</u>, a <u>call expression</u> of an <u>expression statement</u> to a function that returns an <u>impl trait</u> or a <u>dyn trait</u> of that trait violates the <u>unused\_must\_use</u> lint.

```
#[must_use]
trait Critical {}
impl Critical for i32 {}
fn get_critical() -> impl Critical {
    4i32
}
// Violates the `unused_must_use` lint.
get_critical();
```

r[attributes.diagnostics.must\_use.trait-function] When used on a function in a trait declaration, then the behavior also applies when the call expression is a function from an implementation of the trait.

```
trait Trait {
    #[must_use]
    fn use_me(&self) -> i32;
}
impl Trait for i32 {
    fn use_me(&self) -> i32 { 0i32 }
}
// Violates the `unused_must_use` lint.
5i32.use_me();
```

r[attributes.diagnostics.must\_use.trait-impl-function] When used on a function in a trait implementation, the attribute does nothing.

[!NOTE] Trivial no-op expressions containing the value will not violate the lint. Examples include wrapping the value in a type that does not implement <u>Drop</u> and then not using that type and being the final expression of a <u>block expression</u> that is not used.

```
#[must_use]
fn five() -> i32 { 5i32 }
// None of these violate the unused_must_use lint.
(five(),);
Some(five());
{ five() };
if true { five() } else { 0i32 };
match true {
    _ => five()
};
```

[!NOTE] It is idiomatic to use a <u>let statement</u> with a pattern of \_\_\_\_\_\_ when a must-used value is purposely discarded.

```
#[must_use]
fn five() -> i32 { 5i32 }
// Does not violate the unused_must_use lint.
let _ = five();
```

r[attributes.diagnostic.namespace]

#### The diagnostic tool attribute namespace

r[attributes.diagnostic.namespace.intro] The *#*[diagnostic] attribute namespace is a home for attributes to influence compile-time error messages. The hints provided by these attributes are not guaranteed to be used.

r[attributes.diagnostic.namespace.unknown-invalid-syntax] Unknown attributes in this namespace are accepted, though they may emit warnings for unused attributes. Additionally, invalid inputs to known attributes will typically be a warning (see the attribute definitions for details). This is meant to allow adding or discarding attributes and changing inputs in the future to allow changes without the need to keep the non-meaningful attributes or options working.

r[attributes.diagnostic.on\_unimplemented]

#### The diagnostic::on\_unimplemented attribute

r[attributes.diagnostic.on\_unimplemented.intro] The #
[diagnostic::on\_unimplemented] attribute is a hint to the compiler to
supplement the error message that would normally be generated in
scenarios where a trait is required but not implemented on a type.

r[attributes.diagnostic.on\_unimplemented.allowed-positions] The attribute should be placed on a <u>trait declaration</u>, though it is not an error to be located in other positions.

r[attributes.diagnostic.on\_unimplemented.syntax] The attribute uses the [MetaListNameValueStr] syntax to specify its inputs, though any malformed input to the attribute is not considered as an error to provide both forwards and backwards compatibility.

r[attributes.diagnostic.on\_unimplemented.keys] The following keys have the given meaning:

- message --- The text for the top level error message.
- label --- The text for the label shown inline in the broken code in the error message.
- note --- Provides additional notes.

r[attributes.diagnostic.on\_unimplemented.note-repetition] The note option can appear several times, which results in several note messages being emitted.

r[attributes.diagnostic.on\_unimplemented.repetition] If any of the other options appears several times the first occurrence of the relevant option specifies the actually used value. Subsequent occurrences generates a warning.

r[attributes.diagnostic.on\_unimplemented.unknown-keys] A warning is generated for any unknown keys.

r[attributes.diagnostic.on\_unimplemented.format-string] All three options accept a string as an argument, interpreted using the same formatting as a [std::fmt] string.

r[attributes.diagnostic.on\_unimplemented.format-parameters] Format parameters with the given named parameter will be replaced with the following text:

- {Self} --- The name of the type implementing the trait.
- { *GenericParameterName* } --- The name of the generic argument's type for the given generic parameter.

r[attributes.diagnostic.on\_unimplemented.invalid-formats] Any other format parameter will generate a warning, but will otherwise be included in the string as-is.

r[attributes.diagnostic.on\_unimplemented.invalid-string] Invalid format strings may generate a warning, but are otherwise allowed, but may not display as intended. Format specifiers may generate a warning, but are otherwise ignored.

In this example:

```
#[diagnostic::on_unimplemented(
```

```
message = "My Message for `ImportantTrait<{A}>` implemented
for `{Self}`",
    label = "My Label",
    note = "Note 1",
    note = "Note 2"
)]
```

```
trait ImportantTrait<A> {}
fn use_my_trait(_: impl ImportantTrait<i32>) {}
fn main() {
    use_my_trait(String::new());
}
  the compiler may generate an error message which looks like this:
error[E0277]: My Message for `ImportantTrait<i32>` implemented
for `String`
  --> src/main.rs:14:18
14 |
        use_my_trait(String::new());
         ----- ^^^^ My Label
   required by a bound introduced by this call
   = help: the trait `ImportantTrait<i32>` is not implemented
for `String`
   = note: Note 1
   = note: Note 2
```

### r[attributes.diagnostic.do\_not\_recommend]

#### The diagnostic::do\_not\_recommend attribute

r[attributes.diagnostic.do\_not\_recommend.intro] The #
[diagnostic::do\_not\_recommend] attribute is a hint to the compiler to not
show the annotated trait implementation as part of a diagnostic message.

[!NOTE] Suppressing the recommendation can be useful if you know that the recommendation would normally not be useful to the programmer. This often occurs with broad, blanket impls. The recommendation may send the programmer down the wrong path, or the trait implementation may be an internal detail that you don't want to expose, or the bounds may not be able to be satisfied by the programmer.

For example, in an error message about a type not implementing a required trait, the compiler may find a trait implementation that would satisfy the requirements if it weren't for specific bounds in the trait implementation. The compiler may tell the user that there is an impl, but the problem is the bounds in the trait implementation. The **#** [diagnostic::do\_not\_recommend] attribute can be used to tell the compiler to *not* tell the user about the trait implementation, and instead simply tell the user the type doesn't implement the required trait.

r[attributes.diagnostic.do\_not\_recommend.allowed-positions] The attribute should be placed on a <u>trait implementation item</u>, though it is not an error to be located in other positions.

r[attributes.diagnostic.do\_not\_recommend.syntax] The attribute does not accept any arguments, though unexpected arguments are not considered as an error.

In the following example, there is a trait called AsExpression which is used for casting arbitrary types to the Expression type used in an SQL library. There is a method called check which takes an AsExpression.

```
# pub trait Expression {
#
      type SqlType;
# }
#
# pub trait AsExpression<ST> {
      type Expression: Expression<SqlType = ST>;
#
# }
#
# pub struct Text;
# pub struct Integer;
#
# pub struct Bound<T>(T);
# pub struct SelectInt;
#
# impl Expression for SelectInt {
      type SqlType = Integer;
#
# }
```

```
#
# impl<T> Expression for Bound<T> {
#
      type SqlType = T;
# }
#
# impl AsExpression<Integer> for i32 {
#
      type Expression = Bound<Integer>;
# }
#
# impl AsExpression<Text> for &'_ str {
      type Expression = Bound<Text>;
#
# }
#
# impl<T> Foo for T where T: Expression {}
// Uncomment this line to change the recommendation.
// #[diagnostic::do_not_recommend]
impl<T, ST> AsExpression<ST> for T
where
    T: Expression<SqlType = ST>,
{
    type Expression = T;
}
trait Foo: Expression + Sized {
     fn check<T>(&self, _: T) -> <T as AsExpression<<Self as</pre>
Expression>::SqlType>>::Expression
    where
        T: AsExpression<Self::SqlType>,
    {
        todo!()
    }
}
fn main() {
```

```
SelectInt.check("bar");
```

}

The SelectInt type's check method is expecting an Integer type. Calling it with an i32 type works, as it gets converted to an Integer by the AsExpression trait. However, calling it with a string does not, and generates a an error that may look like this:

```
error[E0277]: the trait bound `&str: Expression`
                                                   is
                                                       not
satisfied
 --> src/main.rs:53:15
  53 |
        SelectInt.check("bar");
                        ^^^^ the trait `Expression` is not
implemented for `&str`
        help: the following other types
     =
                                           implement
                                                     trait
`Expression`:
           Bound<T>
           SelectInt
note: required for `&str` to implement `AsExpression<Integer>`
 --> src/main.rs:45:13
  45 | impl<T, ST> AsExpression<ST> for T
               ^^^^^
  Λ
46 | where
47 |
        T: Expression<SqlType = ST>,
             ----- unsatisfied trait bound
```

```
introduced here
```

By adding the #[diagnostic::do\_not\_recommend] attribute to the blanket impl for AsExpression, the message changes to:

```
error[E0277]: the trait bound `&str: AsExpression<Integer>` is
not satisfied
```

```
--> src/main.rs:53:15
|
```

53 | SelectInt.check("bar");

| ^^^^^ the trait `AsExpression<Integer>` is
not implemented for `&str`

= help: the trait `AsExpression<Integer>` is not implemented
for `&str`

T

but trait `AsExpression<Text>` is implemented for it = help: for that trait implementation, expected `Text`, found `Integer`

The first error message includes a somewhat confusing error message about the relationship of &str and Expression, as well as the unsatisfied trait bound in the blanket impl. After adding # [diagnostic::do\_not\_recommend], it no longer considers the blanket impl for the recommendation. The message should be a little clearer, with an indication that a string cannot be converted to an Integer. r[attributes.codegen]

# **Code generation attributes**

The following <u>attributes</u> are used for controlling code generation. r[attributes.codegen.hint]

## **Optimization hints**

r[attributes.codegen.hint.cold-inline] The cold and inline attributes give suggestions to generate code in a way that may be faster than what it would do without the hint. The attributes are only hints, and may be ignored.

r[attributes.codegen.hint.usage] Both attributes can be used on <u>functions</u>. When applied to a function in a <u>trait</u>, they apply only to that function when used as a default function for a trait implementation and not to all trait implementations. The attributes have no effect on a trait function without a body.

r[attributes.codegen.inline]

### The inline attribute

r[attributes.codegen.inline.intro] The *inline attribute* suggests that a copy of the attributed function should be placed in the caller, rather than generating code to call the function where it is defined.

[!NOTE] The **rustc** compiler automatically inlines functions based on internal heuristics. Incorrectly inlining functions can make the program slower, so this attribute should be used with care.

r[attributes.codegen.inline.modes] There are three ways to use the inline attribute:

- *#[inline] suggests* performing an inline expansion.
- #[inline(always)] suggests that an inline expansion should always
  be performed.
- #[inline(never)] suggests that an inline expansion should never be
  performed.

[!NOTE] **#[inline]** in every form is a hint, with no *requirements* on the language to place a copy of the attributed function in the caller.

r[attributes.codegen.cold]

# The cold attribute

The *cold* <u>attribute</u> suggests that the attributed function is unlikely to be called.

r[attributes.codegen.naked]

#### The naked attribute

r[attributes.codegen.naked.intro] The *naked* <u>attribute</u> prevents the compiler from emitting a function prologue and epilogue for the attributed function.

r[attributes.codegen.naked.body] The <u>function body</u> must consist of exactly one <u>naked asm!</u> macro invocation.

r[attributes.codegen.naked.prologue-epilogue] No function prologue or epilogue is generated for the attributed function. The assembly code in the naked\_asm! block constitutes the full body of a naked function.

r[attributes.codegen.naked.unsafe-attribute] The naked attribute is an <u>unsafe attribute</u>. Annotating a function with **#**[unsafe(naked)] comes with the safety obligation that the body must respect the function's calling convention, uphold its signature, and either return or diverge (i.e., not fall through past the end of the assembly code).

r[attributes.codegen.naked.call-stack] The assembly code may assume that the call stack and register state are valid on entry as per the signature and calling convention of the function.

r[attributes.codegen.naked.no-duplication] The assembly code may not be duplicated by the compiler except when monomorphizing polymorphic functions.

[!NOTE] Guaranteeing when the assembly code may or may not be duplicated is important for naked functions that define symbols.

r[attributes.codegen.naked.unused-variables] The <u>unused variables</u> lint is suppressed within naked functions.

r[attributes.codegen.naked.inline] The <u>inline</u> attribute cannot by applied to a naked function.

r[attributes.codegen.naked.track\_caller] The <u>track caller</u> attribute cannot be applied to a naked function.

r[attributes.codegen.naked.testing] The <u>testing attributes</u> cannot be applied to a naked function.

r[attributes.codegen.no\_builtins]

#### The no\_builtins attribute

The *no\_builtins attribute* may be applied at the crate level to disable optimizing certain code patterns to invocations of library functions that are assumed to exist.

r[attributes.codegen.target\_feature]

#### The target\_feature attribute

r[attributes.codegen.target\_feature.intro] The *target\_feature attribute* may be applied to a function to enable code generation of that function for specific platform architecture features. It uses the [MetaListNameValueStr] syntax with a single key of enable whose value is a string of comma-separated feature names to enable.

```
# #[cfg(target_feature = "avx2")]
#[target_feature(enable = "avx2")]
fn foo_avx2() {}
```

r[attributes.codegen.target\_feature.arch] Each <u>target architecture</u> has a set of features that may be enabled. It is an error to specify a feature for a target architecture that the crate is not being compiled for.

r[attributes.codegen.target\_feature.closures] Closures defined within a target\_feature - annotated function inherit the attribute from the enclosing function.

r[attributes.codegen.target\_feature.target-ub] It is <u>undefined behavior</u> to call a function that is compiled with a feature that is not supported on the current platform the code is running on, *except* if the platform explicitly documents this to be safe.

r[attributes.codegen.target\_feature.safety-restrictions] The following restrictions apply unless otherwise specified by the platform rules below:

- Safe #[target\_feature] functions (and closures that inherit the attribute) can only be safely called within a caller that enables all the target\_features that the callee enables. This restriction does not apply in an unsafe context.
- Safe #[target\_feature] functions (and closures that inherit the attribute) can only be coerced to *safe* function pointers in contexts that enable all the target\_features that the coercee enables. This restriction does not apply to unsafe function pointers.

Implicitly enabled features are included in this rule. For example an sse2 function can call ones marked with sse.

```
# #[cfg(target_feature = "sse2")] {
#[target_feature(enable = "sse")]
fn foo_sse() {}
fn bar() {
     // Calling `foo_sse` here is unsafe, as we must ensure
that SSE is
    // available first, even if `sse` is enabled by default on
the target
    // platform or manually enabled as compiler flags.
    unsafe {
        foo_sse();
    }
}
#[target_feature(enable = "sse")]
fn bar_sse() {
    // Calling `foo_sse` here is safe.
    foo_sse();
    || foo_sse();
}
#[target_feature(enable = "sse2")]
fn bar_sse2() {
     // Calling `foo_sse` here is safe because `sse2` implies
`sse`.
    foo_sse();
}
# }
```

r[attributes.codegen.target\_feature.fn-traits] A function with a # [target\_feature] attribute *never* implements the Fn family of traits, although closures inheriting features from the enclosing function do.

r[attributes.codegen.target\_feature.allowed-positions] The #
[target\_feature] attribute is not allowed on the following places:

- [the main function][crate.main]
- a [panic\_handler function][panic.panic\_handler]
- safe trait methods
- safe default functions in traits

r[attributes.codegen.target\_feature.inline] Functions marked with target\_feature are not inlined into a context that does not support the given features. The #[inline(always)] attribute may not be used with a target\_feature attribute.

r[attributes.codegen.target\_feature.availability]

### Available features

The following is a list of the available feature names.

r[attributes.codegen.target\_feature.x86]

#### x86 or x86\_64

Executing code with unsupported features is undefined behavior on this platform. Hence on this platform usage of <code>#[target\_feature]</code> functions follows the [above restrictions][attributes.codegen.target\_feature.safety-restrictions].

Feature	Implicitly Enables	Description
adx		ADX Multi-Precision Add-Carry Instruction Extensions
aes	sse2	AES Advanced Encryption Standard
avx	sse4.2	AVX Advanced Vector Extensions
avx2	avx	AVX2 Advanced Vector Extensions 2
avx51 2bf16	avx512b w	<u>AVX512-BF16</u> Advanced Vector Extensions 512-bit - Bfloat16 Extensions
avx512 bitalg	avx512b W	<u>AVX512-BITALG</u> Advanced Vector Extensions 512-bit - Bit Algorithms

Feature	Implicitly Enables	Description
avx51 2bw	avx512f	<u>AVX512-BW</u> Advanced Vector Extensions 512-bit - Byte and Word Instructions
avx51 2cd	avx512f	<u>AVX512-CD</u> Advanced Vector Extensions 512-bit - Conflict Detection Instructions
avx51 2dq	avx512f	<u>AVX512-DQ</u> Advanced Vector Extensions 512-bit - Doubleword and Quadword Instructions
avx51 2f	avx2, fma, f16c	<u>AVX512-F</u> Advanced Vector Extensions 512-bit - Foundation
avx51 2fp16	avx512b W	<u>AVX512-FP16</u> Advanced Vector Extensions 512-bit - Float16 Extensions
avx51 2ifma	avx512f	<u>AVX512-IFMA</u> Advanced Vector Extensions 512-bit - Integer Fused Multiply Add
avx51 2vbmi	avx512b w	AVX512-VBMI Advanced Vector Extensions 512-bit - Vector Byte Manipulation Instructions
avx512 vbmi2	avx512b W	AVX512-VBMI2 Advanced Vector Extensions 512-bit - Vector Byte Manipulation Instructions 2
avx51 2vl	avx512f	AVX512-VL Advanced Vector Extensions 512-bit - Vector Length Extensions
avx51 2vnni	avx512f	<u>AVX512-VNNI</u> Advanced Vector Extensions 512-bit - Vector Neural Network Instructions

Feature	Implicitly Enables	Description
avx512 vp2int ersect	avx512f	<u>AVX512-VP2INTERSECT</u> Advanced Vector Extensions 512-bit - Vector Pair Intersection to a Pair of Mask Registers
avx512 vpopcn tdq	avx512f	AVX512-VPOPCNTDQ Advanced Vector Extensions 512-bit - Vector Population Count Instruction
avxif ma	avx2	<u>AVX-IFMA</u> Advanced Vector Extensions - Integer Fused Multiply Add
avxnec onvert	avx2	AVX-NE-CONVERT Advanced Vector Extensions - No-Exception Floating-Point conversion Instructions
avxvn ni	avx2	<u>AVX-VNNI</u> Advanced Vector Extensions - Vector Neural Network Instructions
avxvnn iint16	avx2	<u>AVX-VNNI-INT16</u> Advanced Vector Extensions - Vector Neural Network Instructions with 16-bit Integers
avxvnn iint8	avx2	<u>AVX-VNNI-INT8</u> Advanced Vector Extensions - Vector Neural Network Instructions with 8-bit Integers
bmi1		BMI1 Bit Manipulation Instruction Sets
bmi2		<u>BMI2</u> Bit Manipulation Instruction Sets 2
cmpxc hg16b		<u>cmpxchg16b</u> Compares and exchange 16 bytes (128 bits) of data atomically
f16c	avx	<u>F16C</u> 16-bit floating point conversion instructions

Feature	Implicitly Enables	Description
fma	avx	<u>FMA3</u> Three-operand fused multiply- add
fxsr		<b>fxsave</b> and <b>fxrstor</b> Save and restore x87 FPU, MMX Technology, and SSE State
gfni	sse2	GFNI Galois Field New Instructions
kl	sse2	<u>KEYLOCKER</u> Intel Key Locker Instructions
lzcnt		<pre>lzcnt Leading zeros count</pre>
movbe		<u>movbe</u> Move data after swapping bytes
pclmu lqdq	sse2	<pre>pclmulqdq Packed carry-less multiplication quadword</pre>
popcn t		<b>popcnt</b> Count of bits set to 1
rdran d		rdrand Read random number
rdsee d		rdseed Read random seed
sha	sse2	SHA Secure Hash Algorithm
sha51 2	avx2	<u>SHA512</u> Secure Hash Algorithm with 512-bit digest
sm3	avx	<u>SM3</u> ShangMi 3 Hash Algorithm
sm4	avx2	<u>SM4</u> ShangMi 4 Cipher Algorithm
sse		SSE Streaming SIMD Extensions
sse2	sse	SSE2 Streaming SIMD Extensions 2

Feature	Implicitly Enables	Description
sse3	sse2	SSE3 Streaming SIMD Extensions 3
sse4. 1	ssse3	<u>SSE4.1</u> Streaming SIMD Extensions 4.1
sse4. 2	sse4.1	<u>SSE4.2</u> Streaming SIMD Extensions 4.2
ssse3	sse3	SSSE3 Supplemental Streaming SIMD Extensions 3
vaes	avx2, aes	<u>VAES</u> Vector AES Instructions
vpclm ulqdq	avx, pclmulq dq	<b><u>VPCLMULQDQ</u></b> Vector Carry-less multiplication of Quadwords
widek l	kl	<u>KEYLOCKER WIDE</u> Intel Wide Keylocker Instructions
xsave		xsave Save processor extended states
xsave c		<pre>xsavec Save processor extended states with compaction</pre>
xsave opt		<pre>xsaveopt Save processor extended states optimized</pre>
xsave s		xsaves Save processor extended states supervisor

r[attributes.codegen.target\_feature.aarch64]

#### aarch64

On this platform the usage of #[target\_feature] functions follows the [above restrictions][attributes.codegen.target\_feature.safety-restrictions]. Further documentation on these features can be found in the <u>ARM</u> <u>Architecture Reference Manual</u>, or elsewhere on <u>developer.arm.com</u>.

[!NOTE] The following pairs of features should both be marked as enabled or disabled together if used:

• paca and pacg, which LLVM currently implements as one feature.

Feature	Implicitly Enables	Feature Name
aes	neon	FEAT_AES & FEAT_PMULL Advanced <u>SIMD</u> AES & PMULL instructions
bf16		FEAT_BF16 BFloat16 instructions
bti		FEAT_BTI Branch Target Identification
crc		FEAT_CRC CRC32 checksum instructions
dit		FEAT_DIT Data Independent Timing instructions
dotpr od	neon	FEAT_DotProd Advanced SIMD Int8 dot product instructions
dpb		FEAT_DPB Data cache clean to point of persistence
dpb2	dpb	FEAT_DPB2 Data cache clean to point of deep persistence
f32mm	sve	FEAT_F32MM SVE single-precision FP matrix multiply instruction
f64mm	sve	FEAT_F64MM SVE double-precision FP matrix multiply instruction

Feature	Implicitly Enables	Feature Name
fcma	neon	FEAT_FCMA Floating point complex number support
fhm	fp16	FEAT_FHM Half-precision FP FMLAL instructions
flagm		FEAT_FLAGM Conditional flag manipulation
fp16	neon	FEAT_FP16 Half-precision FP data processing
frint ts		FEAT_FRINTTS Floating-point to int helper instructions
i8mm		FEAT_I8MM Int8 Matrix Multiplication
jscon v	neon	FEAT_JSCVT JavaScript conversion instruction
lor		FEAT_LOR Limited Ordering Regions extension
lse		FEAT_LSE Large System Extensions
mte		FEAT_MTE & FEAT_MTE2 Memory Tagging Extension
neon		FEAT_AdvSimd & FEAT_FP Floating Point and Advanced SIMD extension
paca		FEAT_PAUTH Pointer Authentication (address authentication)
pacg		FEAT_PAUTH Pointer Authentication (generic authentication)
pan		FEAT_PAN Privileged Access-Never extension

Feature	Implicitly Enables	Feature Name
pmuv3		FEAT_PMUv3 Performance Monitors extension (v3)
rand		FEAT_RNG Random Number Generator
ras		FEAT_RAS&FEAT_RASv1p1Reliability,AvailabilityandServiceability extension
rcpc		FEAT_LRCPC Release consistent Processor Consistent
rcpc2	rcpc	FEAT_LRCPC2 RcPc with immediate offsets
rdm	neon	FEAT_RDM Rounding Double Multiply accumulate
sb		FEAT_SB Speculation Barrier
sha2	neon	FEAT_SHA1 & FEAT_SHA256 Advanced SIMD SHA instructions
sha3	sha2	FEAT_SHA512 & FEAT_SHA3 Advanced SIMD SHA instructions
sm4	neon	FEAT_SM3 & FEAT_SM4 Advanced SIMD SM3/4 instructions
spe		FEAT_SPE Statistical Profiling Extension
ssbs		FEAT_SSBS & FEAT_SSBS2 Speculative Store Bypass Safe
sve	neon	FEAT_SVE Scalable Vector Extension
sve2	sve	FEAT_SVE2 Scalable Vector Extension 2

Feature	Implicitly Enables	Feature Name
sve2-	sve2,	FEAT_SVE_AES &
aes	aes	FEAT_SVE_PMULL128 SVE AES instructions
sve2-	sve2	FEAT_SVE2_BitPerm SVE Bit
bitper		Permute
m		
sve2-	sve2,	FEAT_SVE2_SHA3 SVE SHA3
sha3	sha3	instructions
sve2-	sve2,	FEAT_SVE2_SM4 SVE SM4
sm4	sm4	instructions
tme		FEAT_TME Transactional Memory
		Extension
vh		FEAT_VHE Virtualization Host
		Extensions

r[attributes.codegen.target\_feature.loongarch]

#### loongarch

On this platform the usage of #[target\_feature] functions follows the [above restrictions][attributes.codegen.target\_feature.safety-restrictions].

Feature	Implicitly Enables	Description
f		<u><b>F</b></u> Single-precision float-point instructions
d	f	<u><b>D</b></u> Double-precision float-point instructions
freci pe		FRECIPEReciprocalapproximation instructions

Feature	Implicitly Enables	Description
lasx	lsx	LASX 256-bit vector instructions
lbt		<u>LBT</u> Binary translation instructions
lsx	d	LSX 128-bit vector instructions
lvz		LVZ Virtualization instructions

r[attributes.codegen.target\_feature.riscv]

#### riscv32 or riscv64

On this platform the usage of #[target\_feature] functions follows the [above restrictions][attributes.codegen.target\_feature.safety-restrictions].

Further documentation on these features can be found in their respective specification. Many specifications are described in the <u>RISC-V ISA Manual</u> or in another manual hosted on the <u>RISC-V GitHub Account</u>.

Feature	Implicitly Enables	Description
a		<u>A</u> Atomic instructions
C		C Compressed instructions
m		<u>M</u> Integer Multiplication and Division instructions
zb	zba, zbc, zbs	<u><b>Zb</b></u> Bit Manipulation instructions
zba		Zba Address Generation instructions
zbb		<u>Zbb</u> Basic bit- manipulation
zbc		Zbc Carry-less multiplication

Feature	Implicitly Enables	Description
zbkb		Zbkb Bit Manipulation Instructions for Cryptography
zbkc		ZbkcCarry-lessmultiplicationforCryptography
zbkx		Zbkx Crossbar permutations
zbs		Zbs Single-bit instructions
zk	zkn, zkr, zks, zkt, zbkb, zbkc, zkbx	Zk Scalar Cryptography
zkn	zknd, zkne, zknh, zbkb, zbkc, zkbx	Zkn NIST Algorithm suite extension
zknd		<u>Zknd</u> NIST Suite: AES Decryption
zkne		<u>Zkne</u> NIST Suite: AES Encryption
zknh		Zknh NIST Suite: Hash Function Instructions
zkr		<u>Zkr</u> Entropy Source Extension
zks	zksed, zksh, zbkb, zbkc, zkbx	<u>Zks</u> ShangMi Algorithm Suite
zksed		Zksed ShangMi Suite: SM4 Block Cipher Instructions
zksh		Zksh ShangMi Suite: SM3 Hash Function Instructions

Feature	Implicitly Enables	Description
zkt		Zkt Data Independent Execution Latency Subset

r[attributes.codegen.target\_feature.wasm]

#### wasm32 or wasm64

Safe **#[target\_feature]** functions may always be used in safe contexts on Wasm platforms. It is impossible to cause undefined behavior via the **#** [target\_feature] attribute because attempting to use instructions unsupported by the Wasm engine will fail at load time without the risk of being interpreted in a way different from what the compiler expected.

Feature	Implicitly Enables	Description
bulk- memory		<u>WebAssembly bulk memory</u> operations proposal
extended- const		<u>WebAssembly extended const</u> <u>expressions proposal</u>
mutable- globals		<u>WebAssembly mutable global</u> <u>proposal</u>
nontrappin g-fptoint		<u>WebAssembly non-trapping float-</u> to-int conversion proposal
relaxed- simd	simd128	<u>WebAssembly relaxed simd</u> <u>proposal</u>
sign-ext		<u>WebAssembly sign extension</u> operators Proposal
simd128		WebAssembly simd proposal
multivalue		<u>WebAssembly multivalue proposal</u>

Feature	Implicitly Enables	Description
reference- types		<u>WebAssembly</u> reference-types proposal
tail-call		WebAssembly tail-call proposal

r[attributes.codegen.target\_feature.info]

### **Additional information**

r[attributes.codegen.target\_feature.remark-cfg] See the <u>target feature</u> <u>conditional compilation option</u> for selectively enabling or disabling compilation of code based on compile-time settings. Note that this option is not affected by the <u>target\_feature</u> attribute, and is only driven by the features enabled for the entire crate.

r[attributes.codegen.target\_feature.remark-rt]Seetheis x86 feature detectedoris aarch64 feature detectedmacros inthe standard library for runtime feature detection on these platforms.

[!NOTE] **rustc** has a default set of features enabled for each target and CPU. The CPU may be chosen with the <u>-C target-cpu</u> flag. Individual features may be enabled or disabled for an entire crate with the <u>-C target-feature</u> flag.

r[attributes.codegen.track\_caller]

### The track\_caller attribute

r[attributes.codegen.track\_caller.allowed-positions] The track\_caller attribute may be applied to any function with <u>"Rust" ABI</u> with the exception of the entry point fn main.

r[attributes.codegen.track\_caller.traits] When applied to functions and methods in trait declarations, the attribute applies to all implementations. If the trait provides a default implementation with the attribute, then the attribute also applies to override implementations.

r[attributes.codegen.track\_caller.extern] When applied to a function in an extern block the attribute must also be applied to any linked implementations, otherwise undefined behavior results. When applied to a function which is made available to an extern block, the declaration in the extern block must also have the attribute, otherwise undefined behavior results.

r[attributes.codegen.track\_caller.behavior]

### **Behavior**

Applying the attribute to a function f allows code within f to get a hint of the Location of the "topmost" tracked call that led to f's invocation. At the point of observation, an implementation behaves as if it walks up the stack from f's frame to find the nearest frame of an *unattributed* function outer, and it returns the Location of the tracked call in outer.

```
#[track_caller]
fn f() {
    println!("{}", std::panic::Location::caller());
}
```

[!NOTE] core provides [core::panic::Location::caller] for observing caller locations. It wraps the [core::intrinsics::caller\_location] intrinsic implemented by rustc. [!NOTE] Because the resulting Location is a hint, an implementation may halt its walk up the stack early. See <u>Limitations</u> for important caveats.

#### Examples

When f is called directly by calls\_f, code in f observes its callsite within calls\_f:

```
# #[track_caller]
# fn f() {
#     println!("{}", std::panic::Location::caller());
# }
fn calls_f() {
     f(); // <-- f() prints this location
}</pre>
```

When f is called by another attributed function g which is in turn called by calls\_g, code in both f and g observes g's callsite within calls\_g:

```
# #[track_caller]
# fn f() {
#     println!("{}", std::panic::Location::caller());
# }
#[track_caller]
fn g() {
     println!("{}", std::panic::Location::caller());
     f();
}
fn calls_g() {
     g(); // <-- g() prints this location twice, once itself
and once from f()
}</pre>
```

When g is called by another attributed function h which is in turn called by calls\_h, all code in f, g, and h observes h's callsite within

```
calls_h:
```

```
# #[track caller]
# fn f() {
      println!("{}", std::panic::Location::caller());
#
# }
# #[track_caller]
# fn g() {
      println!("{}", std::panic::Location::caller());
#
#
      f();
# }
#[track_caller]
fn h() {
    println!("{}", std::panic::Location::caller());
    g();
}
fn calls_h() {
    h(); // <-- prints this location three times, once itself,
once from g(), once from f()
}
```

And so on.

r[attributes.codegen.track\_caller.limits]

### Limitations

r[attributes.codegen.track\_caller.hint] This information is a hint and implementations are not required to preserve it.

r[attributes.codegen.track\_caller.decay] In particular, coercing a function
with #[track\_caller] to a function pointer creates a shim which appears
to observers to have been called at the attributed function's definition site,
losing actual caller information across virtual calls. A common example of
this coercion is the creation of a trait object whose methods are attributed.

[!NOTE] The aforementioned shim for function pointers is necessary because **rustc** implements **track\_caller** in a codegen context by appending an implicit parameter to the function ABI, but
this would be unsound for an indirect call because the parameter is not a part of the function's type and a given function pointer type may or may not refer to a function with the attribute. The creation of a shim hides the implicit parameter from callers of the function pointer, preserving soundness.

r[attributes.codegen.instruction\_set]

## The instruction\_set attribute

r[attributes.codegen.instruction\_set.allowed-positions] The instruction\_set attribute may be applied to a function to control which instruction set the function will be generated for.

r[attributes.codegen.instruction\_set.behavior] This allows mixing more than one instruction set in a single program on CPU architectures that support it.

r[attributes.codegen.instruction\_set.syntax] It uses the [MetaListPaths] syntax, and a path comprised of the architecture family name and instruction set name.

r[attributes.codegen.instruction\_set.target-limits] It is a compilation error to use the instruction\_set attribute on a target that does not support it.

r[attributes.codegen.instruction\_set.arm]

## On ARM

For the ARMv4T and ARMv5te architectures, the following are supported:

```
• arm::a32 --- Generate the function as A32 "ARM" code.
```

```
• arm::t32 --- Generate the function as T32 "Thumb" code.
```

```
#[instruction_set(arm::a32)]
```

```
fn foo_arm_code() {}
```

```
#[instruction_set(arm::t32)]
fn her thumb code() []
```

```
fn bar_thumb_code() {}
```

Using the instruction\_set attribute has the following effects:

- If the address of the function is taken as a function pointer, the low bit of the address will be set to 0 (arm) or 1 (thumb) depending on the instruction set.
- Any inline assembly in the function must use the specified instruction set instead of the target default.

r[attributes.limits]

# Limits

The following <u>attributes</u> affect compile-time limits. r[attributes.limits.recursion\_limit]

### The recursion\_limit attribute

r[attributes.limits.recursion\_limit.intro] The *recursion\_limit* attribute may be applied at the <u>crate</u> level to set the maximum depth for potentially infinitely-recursive compile-time operations like macro expansion or auto-dereference.

r[attributes.limits.recursion\_limit.syntax] It uses the [MetaNameValueStr] syntax to specify the recursion depth.

```
[!NOTE] The default in rustc is 128.
#![recursion limit = "4"]
macro_rules! a {
    () => { a!(1); };
    (1) => { a!(2); };
    (2) => { a!(3); };
    (3) => { a!(4); };
    (4) => \{ \};
}
// This fails to expand because it requires a recursion depth
greater than 4.
a!{}
#![recursion_limit = "1"]
// This fails because it requires two recursive steps to auto-
dereference.
(|_: &u8| {})(&&&1);
```

r[attributes.limits.type\_length\_limit]

#### The type\_length\_limit attribute

[!NOTE] This limit is only enforced when the nightly -Zenforcetype-length-limit flag is active.

For more information, see <u>https://github.com/rust-lang/rust/pull/127670</u>.

r[attributes.limits.type\_length\_limit.intro] The *type\_length\_limit attribute* limits the maximum number of type substitutions made when constructing a concrete type during monomorphization.

r[attributes.limits.type\_length\_limit.syntax] It is applied at the <u>crate</u> level, and uses the [MetaNameValueStr] syntax to set the limit based on the number of type substitutions.

```
[!NOTE] The default in rustc is 1048576.
#![type_length_limit = "4"]
fn f<T>(x: T) {}
// This fails to compile because monomorphizing to
// `f::<((((i32,), i32), i32), i32)>` requires more than 4 type
elements.
f(((((1,), 2), 3), 4));
```

r[attributes.type-system]

# **Type system attributes**

The following <u>attributes</u> are used for changing how a type can be used. r[attributes.type-system.non\_exhaustive]

### The non\_exhaustive attribute

r[attributes.type-system.non\_exhaustive.intro] The *non\_exhaustive attribute* indicates that a type or variant may have more fields or variants added in the future.

r[attributes.type-system.non\_exhaustive.allowed-positions] It can be applied to <u>structs</u>, <u>enums</u>, and <u>enum</u> variants.

r[attributes.type-system.non\_exhaustive.syntax] The non\_exhaustive attribute uses the [MetaWord] syntax and thus does not take any inputs.

r[attributes.type-system.non\_exhaustive.same-crate] Within the defining crate, non\_exhaustive has no effect.

```
#[non_exhaustive]
pub struct Config {
    pub window_width: u16,
    pub window_height: u16,
}
#[non_exhaustive]
pub struct Token;
#[non exhaustive]
pub struct Id(pub u64);
#[non_exhaustive]
pub enum Error {
    Message(String),
    Other,
}
pub enum Message {
     #[non_exhaustive] Send { from: u32, to: u32, contents:
String },
    #[non_exhaustive] Reaction(u32),
    #[non_exhaustive] Quit,
```

}

```
// Non-exhaustive structs can be constructed as normal within
the defining crate.
let config = Config { window_width: 640, window_height: 480 };
let token = Token;
let id = Id(4);
11
   Non-exhaustive structs can be matched on exhaustively
within the defining crate.
let Config { window_width, window_height } = config;
let Token = token;
let Id(id_number) = id;
let error = Error::Other;
let message = Message::Reaction(3);
// Non-exhaustive enums can be matched on exhaustively within
the defining crate.
match error {
    Error::Message(ref s) => { },
    Error::Other => { },
}
match message {
    // Non-exhaustive variants can be matched on exhaustively
within the defining crate.
    Message::Send { from, to, contents } => { },
    Message::Reaction(id) => { },
    Message::Quit => { },
}
```

r[attributes.type-system.non\_exhaustive.external-crate] Outside of the defining crate, types annotated with non\_exhaustive have limitations that preserve backwards compatibility when new fields or variants are added.

r[attributes.type-system.non\_exhaustive.construction] Non-exhaustive types cannot be constructed outside of the defining crate:

- Non-exhaustive variants (<u>struct</u> or <u>enum</u> <u>variant</u>) cannot be constructed with a [StructExpression] (including with <u>functional</u> <u>update syntax</u>).
- The implicitly defined same-named constant of a <u>unit-like struct</u>, or the same-named constructor function of a <u>tuple struct</u>, has a <u>visibility</u> no greater than pub(crate). That is, if the struct's visibility is pub, then the constant or constructor's visibility is pub(crate), and otherwise the visibility of the two items is the same (as is the case without # [non\_exhaustive]).
- <u>enum</u> instances can be constructed.

The following examples of construction do not compile when outside the defining crate:

```
// These are types defined in an upstream crate that have been
annotated as
// `#[non_exhaustive]`.
use upstream::{Config, Token, Id, Error, Message};
// Cannot construct an instance of `Config`; if new fields were
added in
// a new version of `upstream` then this would fail to compile,
so it is
// disallowed.
let config = Config { window_width: 640, window_height: 480 };
// Cannot construct an instance of `Token`; if new fields were
added, then
// it would not be a unit-like struct any more, so the same-
named constant
// created by it being a unit-like struct is not public outside
the crate;
// this code fails to compile.
let token = Token;
```

```
// Cannot construct an instance of `Id`; if new fields were
added, then
// its constructor function signature would change, so its
constructor
// function is not public outside the crate; this code fails to
compile.
let id = Id(5);
// Can construct an instance of `Error`; new variants being
introduced would
// not result in this failing to compile.
let error = Error::Message("foo".to_string());
//
    Cannot
            construct
                            instance
                                          `Message::Send`
                       an
                                      of
                                                           or
`Message::Reaction`;
// if new fields were added in a new version of `upstream` then
this would
// fail to compile, so it is disallowed.
let message = Message::Send { from: 0, to: 1, contents:
"foo".to_string(), };
let message = Message::Reaction(0);
// Cannot construct an instance of `Message::Quit`; if this
were converted to
```

// a tuple-variant `upstream` then this would fail to compile. let message = Message::Quit;

r[attributes.type-system.non\_exhaustive.match] There are limitations when matching on non-exhaustive types outside of the defining crate:

• When pattern matching on a non-exhaustive <u>enum</u>, matching on a variant does not contribute towards the exhaustiveness of the arms.

The following examples of matching do not compile when outside the defining crate:

// These are types defined in an upstream crate that have been
annotated as

```
// `#[non_exhaustive]`.
use upstream::{Config, Token, Id, Error, Message};
// Cannot match on a non-exhaustive enum without including a
wildcard arm.
match error {
   Error::Message(ref s) => { },
   Error::Other => { },
   // would compile with: `_ => {},`
}
// Cannot match on a non-exhaustive struct without a wildcard.
if let Ok(Config { window_width, window_height }) = config {
    // would compile with: `..`
```

```
}
```

// Cannot match a non-exhaustive unit-like or tuple struct except by using // braced struct syntax with a wildcard. // This would compile as `let Token { .. } = token;` let Token = token; // This would compile as `let Id { 0: id\_number, .. } = id;` let Id(id\_number) = id;

match message {
 // Cannot match on a non-exhaustive struct enum variant
without including a wildcard.

```
Message::Send { from, to, contents } => { },
```

// Cannot match on a non-exhaustive tuple or unit enum

```
variant.
  Message::Reaction(type) => { },
  Message::Quit => { },
}
```

It's also not allowed to use numeric casts (as) on enums that contain any non-exhaustive variants.

For example, the following enum can be cast because it doesn't contain any non-exhaustive variants:

```
#[non_exhaustive]
pub enum Example {
    First,
    Second
}
```

However, if the enum contains even a single non-exhaustive variant, casting will result in an error. Consider this modified version of the same enum:

```
#[non_exhaustive]
pub enum EnumWithNonExhaustiveVariants {
    First,
    #[non_exhaustive]
    Second
```

}

use othercrate::EnumWithNonExhaustiveVariants;

// Error: cannot cast an enum with a non-exhaustive variant
when it's defined in another crate

```
let _ = EnumWithNonExhaustiveVariants::First as u8;
```

Non-exhaustive types are always considered inhabited in downstream crates.

r[attributes.debugger]

# **Debugger** attributes

The following <u>attributes</u> are used for enhancing the debugging experience when using third-party debuggers like GDB or WinDbg.

r[attributes.debugger.debugger\_visualizer]

### The debugger\_visualizer attribute

r[attributes.debugger.debugger\_visualizer.intro] The *debugger\_visualizer* attribute can be used to embed a debugger visualizer file into the debug information. This enables an improved debugger experience for displaying values in the debugger.

r[attributes.debugger.debugger\_visualizer.syntax] It uses the [MetaListNameValueStr] syntax to specify its inputs, and must be specified as a crate attribute.

r[attributes.debugger.debugger\_visualizer.natvis]

### Using debugger\_visualizer with Natvis

r[attributes.debugger\_debugger\_visualizer.natvis.intro] Natvis is an XML-based framework for Microsoft debuggers (such as Visual Studio and WinDbg) that uses declarative rules to customize the display of types. For detailed information on the Natvis format, refer to Microsoft's <u>Natvis</u> <u>documentation</u>.

r[attributes.debugger.debugger\_visualizer.natvis.msvc] This attribute only supports embedding Natvis files on -windows-msvc targets.

r[attributes.debugger.debugger\_visualizer.natvis.path] The path to the Natvis file is specified with the natvis\_file key, which is a path relative to the crate source file:

```
#![debugger_visualizer(natvis_file = "Rectangle.natvis")]
struct FancyRect {
    x: f32,
    y: f32,
    dx: f32,
    dy: f32,
}
fn main() {
    let fancy_rect = FancyRect { x: 10.0, y: 10.0, dx: 5.0,
    dy: 5.0 };
```

```
println!("set breakpoint here");
 }
  and Rectangle.natvis contains:
<?xml version="1.0" encoding="utf-8"?>
<AutoVisualizer
xmlns="http://schemas.microsoft.com/vstudio/debugger/natvis/201
⊙">
    <Type Name="foo::FancyRect">
      <DisplayString>({x}, {y}) + ({dx}, {dy})</DisplayString>
      <Expand>
        <Synthetic Name="LowerLeft">
          <DisplayString>({x}, {y})</DisplayString>
        </Synthetic>
        <Synthetic Name="UpperLeft">
          <DisplayString>({x}, {y + dy})</DisplayString>
        </Synthetic>
        <Synthetic Name="UpperRight">
          <DisplayString>({x + dx}, {y + dy})</DisplayString>
        </Synthetic>
        <Synthetic Name="LowerRight">
          <DisplayString>({x + dx}, {y})</DisplayString>
        </Synthetic>
      </Expand>
    </Type>
</AutoVisualizer>
```

When viewed under WinDbg, the fancy\_rect variable would be shown as follows:

```
> Variables:
> fancy_rect: (10.0, 10.0) + (5.0, 5.0)
> LowerLeft: (10.0, 10.0)
> UpperLeft: (10.0, 15.0)
> UpperRight: (15.0, 15.0)
> LowerRight: (15.0, 10.0)
```

r[attributes.debugger.debugger\_visualizer.gdb]

### Using debugger\_visualizer with GDB

r[attributes.debugger\_debugger\_visualizer.gdb.pretty] GDB supports the use of a structured Python script, called a *pretty printer*, that describes how a type should be visualized in the debugger view. For detailed information on pretty printers, refer to GDB's <u>pretty printing documentation</u>.

Embedded pretty printers are not automatically loaded when debugging a binary under GDB. There are two ways to enable auto-loading embedded pretty printers:

- Launch GDB with extra arguments to explicitly add a directory or binary to the auto-load safe path: gdb -iex "add-auto-load-safepath safe-path path/to/binary" path/to/binary For more information, see GDB's <u>auto-loading documentation</u>.
- 2. Create a file named gdbinit under \$HOME/.config/gdb (you may need to create the directory if it doesn't already exist). Add the following line to that file: add-auto-load-safe-path path/to/binary.

r[attributes.debugger.debugger\_visualizer.gdb.path] These scripts are embedded using the gdb\_script\_file key, which is a path relative to the crate source file.

```
#![debugger_visualizer(gdb_script_file = "printer.py")]
struct Person {
    name: String,
    age: i32,
}
fn main() {
    let bob = Person { name: String::from("Bob"), age: 10 };
    println!("set breakpoint here");
}
```

```
and printer.py contains:
```

```
import gdb
class PersonPrinter:
    "Print a Person"
    def __init__(self, val):
        self.val = val
        self.name = val["name"]
        self.age = int(val["age"])
    def to_string(self):
              return "{} is {} years old.".format(self.name,
self.age)
def lookup(val):
    lookup_tag = val.type.tag
    if lookup_tag is None:
        return None
    if "foo::Person" == lookup_tag:
        return PersonPrinter(val)
    return None
gdb.current_objfile().pretty_printers.append(lookup)
```

When the crate's debug executable is passed into GDB<sup>1</sup>, print bob will display:

```
"Bob" is 10 years old.
1
```

Note: This assumes you are using the rust-gdb script which configures pretty-printers for standard library types like String.

r[attributes.debugger.collapse\_debuginfo]

### The collapse\_debuginfo attribute

r[attributes.debugger.collapse\_debuginfo.intro] The *collapse\_debuginfo attribute* controls whether code locations from a macro definition are collapsed into a single location associated with the macro's call site, when generating debuginfo for code calling this macro.

r[attributes.debugger.collapse\_debuginfo.syntax] The attribute uses the [MetaListIdents] syntax to specify its inputs, and can only be applied to macro definitions.

r[attributes.debugger.collapse\_debuginfo.options] Accepted options:

- #[collapse\_debuginfo(yes)] --- code locations in debuginfo are collapsed.
- #[collapse\_debuginfo(no)] --- code locations in debuginfo are not collapsed.
- #[collapse\_debuginfo(external)] --- code locations in debuginfo
  are collapsed only if the macro comes from a different crate.

r[attributes.debugger.collapse\_debuginfo.default] The external behavior is the default for macros that don't have this attribute, unless they are built-in macros. For built-in macros the default is yes.

```
[!NOTE] rustc has a -C collapse-macro-debuginfo CLI option
to override both the default collapsing behavior and #
[collapse_debuginfo] attributes.
```

```
#[collapse_debuginfo(yes)]
macro_rules! example {
   () => {
      println!("hello!");
   };
}
```

r[stmt-expr]

# **Statements and expressions**

Rust is *primarily* an expression language. This means that most forms of value-producing or effect-causing evaluation are directed by the uniform syntax category of *expressions*. Each kind of expression can typically *nest* within each other kind of expression, and rules for evaluation of expressions involve specifying both the value produced by the expression and the order in which its sub-expressions are themselves evaluated.

In contrast, statements serve *mostly* to contain and explicitly sequence expression evaluation.

r[statement]

## **Statements**

r[statement.syntax]

Statement ->

`;`

| Item

| LetStatement

| ExpressionStatement

| OuterAttribute\* MacroInvocationSemi

r[statement.intro] A *statement* is a component of a <u>block</u>, which is in turn a component of an outer <u>expression</u> or <u>function</u>.

r[statement.kind] Rust has two kinds of statement: <u>declaration</u> <u>statements</u> and <u>expression statements</u>.

r[statement.decl]

### **Declaration statements**

A *declaration statement* is one that introduces one or more *names* into the enclosing statement block. The declared names may denote new variables or new <u>items</u>.

The two kinds of declaration statements are item declarations and let statements.

r[statement.item]

### **Item declarations**

r[statement.item.intro] An *item declaration statement* has a syntactic form identical to an <u>item declaration</u> within a <u>module</u>.

r[statement.item.scope] Declaring an item within a statement block restricts its <u>scope</u> to the block containing the statement. The item is not given a <u>canonical path</u> nor are any sub-items it may declare.

r[statement.item.associated-scope] The exception to this is that associated items defined by <u>implementations</u> are still accessible in outer scopes as long as the item and, if applicable, trait are accessible. It is otherwise identical in meaning to declaring the item inside a module.

r[statement.item.outer-generics] There is no implicit capture of the containing function's generic parameters, parameters, and local variables. For example, inner may not access outer\_var.

```
fn outer() {
   let outer_var = true;
   fn inner() { /* outer_var is not in scope here */ }
   inner();
}
```

r[statement.let]

#### let statements

r[statement.let.syntax]

r[statement.let.intro] A *let statement* introduces a new set of <u>variables</u>, given by a <u>pattern</u>. The pattern is followed optionally by a type annotation and then either ends, or is followed by an initializer expression plus an optional else block.

r[statement.let.inference] When no type annotation is given, the compiler will infer the type, or signal an error if insufficient type information is available for definite inference.

r[statement.let.scope] Any variables introduced by a variable declaration are visible from the point of declaration until the end of the enclosing block scope, except when they are shadowed by another variable declaration.

r[statement.let.constraint] If an else block is not present, the pattern must be irrefutable. If an else block is present, the pattern may be refutable.

r[statement.let.behavior] If the pattern does not match (this requires it to be refutable), the else block is executed. The else block must always diverge (evaluate to the <u>never type</u>).

panic!();

#### };

r[statement.expr]

### **Expression statements**

```
r[statement.expr.syntax]
```

```
ExpressionStatement ->
```

ExpressionWithoutBlock `;`

| ExpressionWithBlock `;`?

r[statement.expr.intro] An *expression statement* is one that evaluates an <u>expression</u> and ignores its result. As a rule, an expression statement's purpose is to trigger the effects of evaluating its expression.

r[statement.expr.restriction-semicolon] An expression that consists of only a <u>block expression</u> or control flow expression, if used in a context where a statement is permitted, can omit the trailing semicolon. This can cause an ambiguity between it being parsed as a standalone statement and as a part of another expression; in this case, it is parsed as a statement.

r[statement.expr.constraint-block] The type of [ExpressionWithBlock] expressions when used as statements must be the unit type.

```
# let mut v = vec![1, 2, 3];
v.pop(); // Ignore the element returned from pop
if v.is_empty() {
    v.push(5);
} else {
    v.remove(0);
} // Semicolon can be omitted.
[1]; // Separate expression statement, not an
indexing expression.
```

When the trailing semicolon is omitted, the result must be type ().

```
// bad: the block's type is i32, not ()
// Error: expected `()` because of default return type
// if true {
// 1
// 3
// good: the block's type is i32
if true {
```

1
} else {
2
};

r[statement.attribute]

### **Attributes on Statements**

Statements accept <u>outer attributes</u>. The attributes that have meaning on a statement are <u>cfg</u>, and <u>the lint check attributes</u>.

r[expr]

# Expressions

r[expr.syntax] Expression -> ExpressionWithoutBlock | ExpressionWithBlock ExpressionWithoutBlock -> OuterAttribute\* ( LiteralExpression | PathExpression | OperatorExpression | GroupedExpression | ArrayExpression | AwaitExpression | IndexExpression | TupleExpression | TupleIndexingExpression | StructExpression | CallExpression | MethodCallExpression | FieldExpression | ClosureExpression | AsyncBlockExpression | ContinueExpression | BreakExpression | RangeExpression | ReturnExpression | UnderscoreExpression | MacroInvocation )

ExpressionWithBlock ->

```
OuterAttribute*
(
BlockExpression
| ConstBlockExpression
| UnsafeBlockExpression
| LoopExpression
| IfExpression
| MatchExpression
)
```

r[expr.intro] An expression may have two roles: it always produces a *value*, and it may have *effects* (otherwise known as "side effects").

r[expr.evaluation] An expression *evaluates to* a value, and has effects during *evaluation*.

r[expr.operands] Many expressions contain sub-expressions, called the *operands* of the expression.

r[expr.behavior] The meaning of each kind of expression dictates several things:

- Whether or not to evaluate the operands when evaluating the expression
- The order in which to evaluate the operands
- How to combine the operands' values to obtain the value of the expression

r[expr.structure] In this way, the structure of expressions dictates the structure of execution. Blocks are just another kind of expression, so blocks, statements, expressions, and blocks again can recursively nest inside each other to an arbitrary depth.

[!NOTE] We give names to the operands of expressions so that we may discuss them, but these names are not stable and may be changed.

r[expr.precedence]

## **Expression precedence**

The precedence of Rust operators and expressions is ordered as follows, going from strong to weak. Binary Operators at the same precedence level are grouped in the order given by their associativity.

Operator/Expression	Associativity
[Paths][expr.path]	
[Method calls][expr.method]	
[Field expressions][expr.field]	left to right
[Function calls][expr.call], [array indexing] [expr.array.index]	
[?][expr.try]	
Unary [ - ][expr.negate] [ ! ][expr.negate] [ * ] [expr.deref] [borrow][expr.operator.borrow]	
[as][expr.as]	left to right
[ * ][expr.arith-logic] [ / ][expr.arith-logic] [ % ] [expr.arith-logic]	left to right
[ + ][expr.arith-logic] [ - ][expr.arith-logic]	left to right
[ << ][expr.arith-logic] [ >> ][expr.arith-logic]	left to right
[&][expr.arith-logic]	left to right
[ ^ ][expr.arith-logic]	left to right
[ ][expr.arith-logic]	left to right
[ == ][expr.cmp] [ != ][expr.cmp] [ < ][expr.cmp] [ > ][expr.cmp] [ <= ][expr.cmp] [ >= ][expr.cmp]	Require parentheses
[ && ][expr.bool-logic]	left to right
[    ][expr.bool-logic]	left to right
[ ][expr.range] [ = ][expr.range]	Require parentheses

Operator/Expression	Associativity
[=][expr.assign] [+=][expr.compound-assign] [-	right to left
= ][expr.compound-assign] [ *= ][expr.compound-	
assign] [/=][expr.compound-assign] [%=]	
[expr.compound-assign]	
[ &= ][expr.compound-assign] [  = ]	
[expr.compound-assign] [ ^= ][expr.compound-	
assign] [ <<= ][expr.compound-assign] [ >>= ]	
[expr.compound-assign]	
<pre>[return][expr.return] [break][expr.loop.break]</pre>	
[closures][expr.closure]	

r[expr.operand-order]
#### **Evaluation order of operands**

r[expr.operand-order.default] The following list of expressions all evaluate their operands the same way, as described after the list. Other expressions either don't take operands or evaluate them conditionally as described on their respective pages.

- Dereference expression
- Error propagation expression
- Negation expression
- Arithmetic and logical binary operators
- Comparison operators
- Type cast expression
- Grouped expression
- Array expression
- Await expression
- Index expression
- Tuple expression
- Tuple index expression
- Struct expression
- Call expression
- Method call expression
- Field expression
- Break expression
- Range expression
- Return expression

r[expr.operand-order.operands-before-primary] The operands of these expressions are evaluated prior to applying the effects of the expression. Expressions taking multiple operands are evaluated left to right as written in the source code.

[!NOTE] Which subexpressions are the operands of an expression is determined by expression precedence as per the previous section.

For example, the two next method calls will always be called in the same order:

```
# // Using vec instead of array to avoid references
# // since there is no stable owned array iterator
# // at the time this example was written.
let mut one_two = vec![1, 2].into_iter();
assert_eq!(
        (1, 2),
        (one_two.next().unwrap(), one_two.next().unwrap())
);
```

[!NOTE] Since this is applied recursively, these expressions are also evaluated from innermost to outermost, ignoring siblings until there are no inner subexpressions.

r[expr.place-value]

#### **Place Expressions and Value Expressions**

r[expr.place-value.intro] Expressions are divided into two main categories: place expressions and value expressions; there is also a third, minor category of expressions called assignee expressions. Within each expression, operands may likewise occur in either place context or value context. The evaluation of an expression depends both on its own category and the context it occurs within.

r[expr.place-value.place-memory-location] A *place expression* is an expression that represents a memory location.

r[expr.place-value.place-expr-kinds] These expressions are <u>paths</u> which refer to local variables, <u>static variables</u>, <u>dereferences</u> (\*expr), <u>array</u> <u>indexing</u> expressions (expr[expr]), <u>field</u> references (expr.f) and parenthesized place expressions.

r[expr.place-value.value-expr-kinds] All other expressions are value expressions.

r[expr.place-value.value-result] A *value expression* is an expression that represents an actual value.

r[expr.place-value.place-context] The following contexts are *place expression* contexts:

- The left operand of a <u>compound assignment</u> expression.
- The operand of a unary <u>borrow</u>, <u>raw borrow</u> or <u>dereference</u> operator.
- The operand of a field expression.
- The indexed operand of an array indexing expression.
- The operand of any <u>implicit borrow</u>.
- The initializer of a <u>let statement</u>.
- The <u>scrutinee</u> of an <u>if let</u>, <u>match</u>, or <u>while let</u> expression.
- The base of a <u>functional update</u> struct expression.

[!NOTE] Historically, place expressions were called *lvalues* and value expressions were called *rvalues*.

r[expr.place-value.assignee] An *assignee expression* is an expression that appears in the left operand of an <u>assignment</u> expression. Explicitly, the

assignee expressions are:

- Place expressions.
- <u>Underscores</u>.
- <u>Tuples</u> of assignee expressions.
- [Slices][expr.array.index] of assignee expressions.
- <u>Tuple structs</u> of assignee expressions.
- <u>Structs</u> of assignee expressions (with optionally named fields).
- <u>Unit structs</u>

r[expr.place-value.parenthesis] Arbitrary parenthesisation is permitted inside assignee expressions.

r[expr.move]

#### Moved and copied types

r[expr.move.intro] When a place expression is evaluated in a value expression context, or is bound by value in a pattern, it denotes the value held *in* that memory location.

r[expr.move.copy] If the type of that value implements <u>Copy</u>, then the value will be copied.

r[expr.move.requires-sized] In the remaining situations, if that type is <u>Sized</u>, then it may be possible to move the value.

r[expr.move.movable-place] Only the following place expressions may be moved out of:

- <u>Variables</u> which are not currently borrowed.
- <u>Temporary values</u>.
- <u>Fields</u> of a place expression which can be moved out of and don't implement <u>Drop</u>.
- The result of <u>dereferencing</u> an expression with type [Box<T>] and that can also be moved out of.

r[expr.move.deinitialization] After moving out of a place expression that evaluates to a local variable, the location is deinitialized and cannot be read from again until it is reinitialized. r[expr.move.place-invalid] In all other cases, trying to use a place expression in a value expression context is an error.

r[expr.mut]

# Mutability

r[expr.mut.intro] For a place expression to be <u>assigned</u> to, mutably <u>borrowed</u>, <u>implicitly mutably borrowed</u>, or bound to a pattern containing ref mut, it must be *mutable*. We call these *mutable place expressions*. In contrast, other place expressions are called *immutable place expressions*.

r[expr.mut.valid-places] The following expressions can be mutable place expression contexts:

- Mutable <u>variables</u> which are not currently borrowed.
- <u>Mutable static items</u>.
- <u>Temporary values</u>.
- <u>Fields</u>: this evaluates the subexpression in a mutable place expression context.
- <u>Dereferences</u> of a \*mut T pointer.
- Dereference of a variable, or field of a variable, with type &mut T. Note: This is an exception to the requirement of the next rule.
- Dereferences of a type that implements DerefMut : this then requires that the value being dereferenced is evaluated in a mutable place expression context.
- <u>Array indexing</u> of a type that implements **IndexMut**: this then evaluates the value being indexed, but not the index, in mutable place expression context.

r[expr.temporary]

# Temporaries

When using a value expression in most place expression contexts, a temporary unnamed memory location is created and initialized to that value. The expression evaluates to that location instead, except if <u>promoted</u> to a **static**. The <u>drop scope</u> of the temporary is usually the end of the enclosing statement.

r[expr.implicit-borrow]

# **Implicit Borrows**

r[expr.implicit-borrow-intro] Certain expressions will treat an expression as a place expression by implicitly borrowing it. For example, it is possible to compare two unsized <u>slices</u> for equality directly, because the == operator implicitly borrows its operands:

```
# let c = [1, 2, 3];
# let d = vec![1, 2, 3];
let a: &[i32];
let b: &[i32];
# a = &c;
# b = &d;
// ...
*a == *b;
// Equivalent form:
::std::cmp::PartialEq::eq(&*a, &*b);
```

r[expr.implicit-borrow.application] Implicit borrows may be taken in the following expressions:

- Left operand in <u>method-call</u> expressions.
- Left operand in <u>field</u> expressions.
- Left operand in <u>call expressions</u>.
- Left operand in <u>array indexing</u> expressions.
- Operand of the <u>dereference operator</u> (\*).
- Operands of <u>comparison</u>.
- Left operands of the <u>compound assignment</u>. r[expr.overload]

#### **Overloading Traits**

Many of the following operators and expressions can also be overloaded for other types using traits in std::ops or std::cmp. These traits also exist in core::ops and core::cmp with the same names.

r[expr.attr]

## **Expression Attributes**

r[expr.attr.restriction] <u>Outer attributes</u> before an expression are allowed only in a few specific cases:

- Before an expression used as a <u>statement</u>.
- Elements of <u>array expressions</u>, <u>tuple expressions</u>, <u>call expressions</u>, and tuple-style <u>struct</u> expressions.
- The tail expression of <u>block expressions</u>.

r[expr.attr.never-before] They are never allowed before:

- <u>Range</u> expressions.
- Binary operator expressions ([ArithmeticOrLogicalExpression], [ComparisonExpression], [LazyBooleanExpression], [TypeCastExpression], [AssignmentExpression], [CompoundAssignmentExpression]).

r[expr.literal]

# Literal expressions

r[expr.literal.syntax]

```
LiteralExpression ->
```

CHAR\_LITERAL

```
| STRING_LITERAL
```

```
| RAW_STRING_LITERAL
```

```
| BYTE_LITERAL
```

```
| BYTE_STRING_LITERAL
```

```
| RAW_BYTE_STRING_LITERAL
```

```
| C_STRING_LITERAL
```

```
| RAW_C_STRING_LITERAL
```

```
| INTEGER_LITERAL
```

```
| FLOAT_LITERAL
```

```
| `true`
```

```
| `false`
```

r[expr.literal.intro] A *literal expression* is an expression consisting of a single token, rather than a sequence of tokens, that immediately and directly denotes the value it evaluates to, rather than referring to it by name or some other evaluation rule.

r[expr.literal.const-expr] A literal is a form of <u>constant expression</u>, so is evaluated (primarily) at compile time.

r[expr.literal.literal-token] Each of the lexical <u>literal</u> forms described earlier can make up a literal expression, as can the keywords true and false.

"hello";	//	string type
'5';	//	character type
5;	//	integer type

r[expr.literal.string-representation] In the descriptions below, the *string representation* of a token is the sequence of characters from the input which matched the token's production in a *Lexer* grammar snippet.

[!NOTE] This string representation never includes a character U+000D (CR) immediately followed by U+000A (LF): this pair would have been previously transformed into a single U+000A (LF).

r[expr.literal.escape]

#### Escapes

r[expr.literal.escape.intro] The descriptions of textual literal expressions below make use of several forms of *escape*.

r[expr.literal.escape.sequence] Each form of escape is characterised by:

- an *escape sequence*: a sequence of characters, which always begins with U+005C (\)
- an *escaped value*: either a single character or an empty sequence of characters

In the definitions of escapes below:

- An *octal digit* is any of the characters in the range [0-7].
- A *hexadecimal digit* is any of the characters in the ranges [0-9], [a-f], or [A-F].

r[expr.literal.escape.simple]

#### **Simple escapes**

Each sequence of characters occurring in the first column of the following table is an escape sequence.

In each case, the escaped value is the character given in the corresponding entry in the second column.

Escape sequence	Escaped value
\0	U+0000 (NUL)
\t	U+0009 (HT)
\n	U+000A (LF)
\r	U+000D (CR)
\"	U+0022 (QUOTATION MARK)
$\mathbf{N}^{\mathbf{r}}$	U+0027 (APOSTROPHE)
	U+005C (REVERSE SOLIDUS)

r[expr.literal.escape.hex-octet]

#### 8-bit escapes

The escape sequence consists of  $\setminus x$  followed by two hexadecimal digits.

The escaped value is the character whose <u>Unicode scalar value</u> is the result of interpreting the final two characters in the escape sequence as a hexadecimal integer, as if by [u8::from\_str\_radix] with radix 16.

[!NOTE] The escaped value therefore has a <u>Unicode scalar value</u> in the range of <u>u8</u>.

r[expr.literal.escape.hex-ascii]

## 7-bit escapes

The escape sequence consists of  $\searrow$  followed by an octal digit then a hexadecimal digit.

The escaped value is the character whose <u>Unicode scalar value</u> is the result of interpreting the final two characters in the escape sequence as a hexadecimal integer, as if by [u8::from\_str\_radix] with radix 16.

r[expr.literal.escape.unicode]

## **Unicode escapes**

The escape sequence consists of \u{, followed by a sequence of characters each of which is a hexadecimal digit or \_, followed by }.

The escaped value is the character whose <u>Unicode scalar value</u> is the result of interpreting the hexadecimal digits contained in the escape sequence as a hexadecimal integer, as if by [u32::from\_str\_radix] with radix 16.

[!NOTE] The permitted forms of a [CHAR\_LITERAL] or [STRING\_LITERAL] token ensure that there is such a character.

r[expr.literal.continuation]

## String continuation escapes

The escape sequence consists of  $\$  followed immediately by U+000A (LF), and all following whitespace characters before the next nonwhitespace character. For this purpose, the whitespace characters are U+0009 (HT), U+000A (LF), U+000D (CR), and U+0020 (SPACE).

The escaped value is an empty sequence of characters.

[!NOTE] The effect of this form of escape is that a string continuation skips following whitespace, including additional newlines. Thus a, b and c are equal:

Skipping additional newlines (as in example c) is potentially confusing and unexpected. This behavior may be adjusted in the future. Until a decision is made, it is recommended to avoid relying on skipping multiple newlines with line continuations. See <u>this issue</u> for more information.

r[expr.literal.char]

## **Character literal expressions**

r[expr.literal.char.intro] A character literal expression consists of a single [CHAR\_LITERAL] token.

r[expr.literal.char.type] The expression's type is the primitive <u>char</u> type.

r[expr.literal.char.no-suffix] The token must not have a suffix.

r[expr.literal.char.literal-content] The token's *literal content* is the sequence of characters following the first U+0027 (') and preceding the last U+0027 (') in the string representation of the token.

r[expr.literal.char.represented] The literal expression's *represented character* is derived from the literal content as follows:

r[expr.literal.char.escape]

- If the literal content is one of the following forms of escape sequence, the represented character is the escape sequence's escaped value:
  - <u>Simple escapes</u>
  - <u>7-bit escapes</u>
  - <u>Unicode escapes</u>

r[expr.literal.char.single]

• Otherwise the represented character is the single character that makes up the literal content.

r[expr.literal.char.result] The expression's value is the <u>char</u> corresponding to the represented character's <u>Unicode scalar value</u>.

[!NOTE] The permitted forms of a [CHAR\_LITERAL] token ensure that these rules always produce a single character.

Examples of character literal expressions:

'R'; // R '\'; // ' '\x52'; // R '\u{00E6}'; // LATIN SMALL LETTER AE (U+00E6) r[expr.literal.string]

## **String literal expressions**

r[expr.literal.string.intro] A string literal expression consists of a single [STRING\_LITERAL] or [RAW\_STRING\_LITERAL] token.

r[expr.literal.string.type] The expression's type is a shared reference
(with static lifetime) to the primitive str type. That is, the type is
&'static str.

r[expr.literal.string.no-suffix] The token must not have a suffix.

r[expr.literal.string.literal-content] The token's *literal content* is the sequence of characters following the first U+0022 (") and preceding the last U+0022 (") in the string representation of the token.

r[expr.literal.string.represented] The literal expression's *represented string* is a sequence of characters derived from the literal content as follows:

r[expr.literal.string.escape]

- If the token is a [STRING\_LITERAL], each escape sequence of any of the following forms occurring in the literal content is replaced by the escape sequence's escaped value.
  - <u>Simple escapes</u>
  - <u>7-bit escapes</u>
  - <u>Unicode escapes</u>
  - <u>String continuation escapes</u>

These replacements take place in left-to-right order. For example, the token " $\$  " $\$  a d t.

r[expr.literal.string.raw]

• If the token is a [RAW\_STRING\_LITERAL], the represented string is identical to the literal content.

r[expr.literal.string.result] The expression's value is a reference to a statically allocated <u>str</u> containing the UTF-8 encoding of the represented string.

Examples of string literal expressions:

"foo"; r"foo";	// foo
"\"foo\"";	// "foo"
"foo #\"# bar";	
r##"foo #"# bar"##;	// foo #"# bar
"\x52"; "R"; r"R";	// R
"\\x52"; r"\x52";	// \x52

r[expr.literal.byte-char]

# **Byte literal expressions**

r[expr.literal.byte-char.intro] A byte literal expression consists of a single [BYTE\_LITERAL] token.

r[expr.literal.byte-char.literal] The expression's type is the primitive u8 type.

r[expr.literal.byte-char.no-suffix] The token must not have a suffix.

r[expr.literal.byte-char.literal-content] The token's *literal content* is the sequence of characters following the first U+0027 (') and preceding the last U+0027 (') in the string representation of the token.

r[expr.literal.byte-char.represented] The literal expression's *represented character* is derived from the literal content as follows:

r[expr.literal.byte-char.escape]

- If the literal content is one of the following forms of escape sequence, the represented character is the escape sequence's escaped value:
  - <u>Simple escapes</u>
  - <u>8-bit escapes</u>

r[expr.literal.byte-char.single]

• Otherwise the represented character is the single character that makes up the literal content.

r[expr.literal.byte-char.result] The expression's value is the represented character's <u>Unicode scalar value</u>.

[!NOTE] The permitted forms of a [BYTE\_LITERAL] token ensure that these rules always produce a single character, whose Unicode scalar value is in the range of <u>u8</u>.

Examples of byte literal expressions:

b'R';	//	82
b'\'';	//	39
b'\x52';	//	82
b'\xA0';	//	160

r[expr.literal.byte-string]

#### **Byte string literal expressions**

r[expr.literal.byte-string.intro] A byte string literal expression consists of a single [BYTE\_STRING\_LITERAL] or [RAW\_BYTE\_STRING\_LITERAL] token.

r[expr.literal.byte-string.type] The expression's type is a shared reference (with static lifetime) to an array whose element type is u8. That is, the type is <code>&'static [u8; N]</code>, where <code>N</code> is the number of bytes in the represented string described below.

r[expr.literal.byte-string.no-suffix] The token must not have a suffix.

r[expr.literal.byte-string.literal-content] The token's *literal content* is the sequence of characters following the first U+0022 (") and preceding the last U+0022 (") in the string representation of the token.

r[expr.literal.byte-string.represented] The literal expression's *represented string* is a sequence of characters derived from the literal content as follows:

r[expr.literal.byte-string.escape]

- If the token is a [BYTE\_STRING\_LITERAL], each escape sequence of any of the following forms occurring in the literal content is replaced by the escape sequence's escaped value.
  - <u>Simple escapes</u>
  - <u>8-bit escapes</u>
  - <u>String continuation escapes</u>

These replacements take place in left-to-right order. For example, the token  $b'' \rightarrow 1$  is converted to the characters  $x \neq 1$ .

r[expr.literal.byte-string.raw]

• If the token is a [RAW\_BYTE\_STRING\_LITERAL], the represented string is identical to the literal content.

r[expr.literal.byte-string.result] The expression's value is a reference to a statically allocated array containing the <u>Unicode scalar values</u> of the characters in the represented string, in the same order.

[!NOTE] The permitted forms of [BYTE\_STRING\_LITERAL] and [RAW\_BYTE\_STRING\_LITERAL] tokens ensure that these rules always produce array element values in the range of <u>u8</u>.

Examples of byte string literal expressions:

b"foo"; br"foo";	//	foo
b"\"foo\"";	//	"foo"
b"foo #\"# bar";		
br##"foo #"# bar"##;	//	foo #"# bar
b"\x52";	//	R
b"\\x52"; br"\x52";	//	\x52

r[expr.literal.c-string]

### **C** string literal expressions

r[expr.literal.c-string.intro] A C string literal expression consists of a single [C\_STRING\_LITERAL] or [RAW\_C\_STRING\_LITERAL] token.

r[expr.literal.c-string.type] The expression's type is a shared reference
(with static lifetime) to the standard library CStr type. That is, the type is
&'static core::ffi::CStr.

r[expr.literal.c-string.no-suffix] The token must not have a suffix.

r[expr.literal.c-string.literal-content] The token's *literal content* is the sequence of characters following the first " and preceding the last " in the string representation of the token.

r[expr.literal.c-string.represented] The literal expression's *represented bytes* are a sequence of bytes derived from the literal content as follows:

r[expr.literal.c-string.escape]

- If the token is a [C\_STRING\_LITERAL], the literal content is treated as a sequence of items, each of which is either a single Unicode character other than \or or an <u>escape</u>. The sequence of items is converted to a sequence of bytes as follows:
  - Each single Unicode character contributes its UTF-8 representation.
  - Each <u>simple escape</u> contributes the <u>Unicode scalar value</u> of its escaped value.
  - Each <u>8-bit escape</u> contributes a single byte containing the <u>Unicode scalar value</u> of its escaped value.
  - Each <u>unicode escape</u> contributes the UTF-8 representation of its escaped value.
  - Each <u>string continuation escape</u> contributes no bytes.

r[expr.literal.c-string.raw]

• If the token is a [RAW\_C\_STRING\_LITERAL], the represented bytes are the UTF-8 encoding of the literal content.

[!NOTE] The permitted forms of [C\_STRING\_LITERAL] and [RAW\_C\_STRING\_LITERAL] tokens ensure that the represented bytes never include a null byte.

r[expr.literal.c-string.result] The expression's value is a reference to a statically allocated <u>CStr</u> whose array of bytes contains the represented bytes followed by a null byte.

Examples of C string literal expressions:

c"foo"; cr"foo";	// foo
c"\"foo\"";	// "foo"
c"foo #\"# bar";	
cr##"foo #"# bar"##;	// foo #"# bar
c"\x52"; c"R"; cr"R";	// R
c"\\x52"; cr"\x52";	// \x52
c"æ";	// LATIN SMALL LETTER AE
(U+00E6)	
c"\u{00E6}";	// LATIN SMALL LETTER AE
(U+00E6)	
c"\xC3\xA6";	// LATIN SMALL LETTER AE
(U+00E6)	
c"\xE6".to_bytes();	// [230]
c"\u{00E6}".to_bytes();	// [195, 166]

r[expr.literal.int]

### **Integer literal expressions**

r[expr.literal.int.intro] An integer literal expression consists of a single [INTEGER\_LITERAL] token.

r[expr.literal.int.suffix] If the token has a <u>suffix</u>, the suffix must be the name of one of the <u>primitive integer types</u>: u8, i8, u16, i16, u32, i32, u64, i64, u128, i128, usize, or isize, and the expression has that type.

r[expr.literal.int.infer] If the token has no suffix, the expression's type is determined by type inference:

r[expr.literal.int.inference-unique-type]

• If an integer type can be *uniquely* determined from the surrounding program context, the expression has that type.

r[expr.literal.int.inference-default]

• If the program context under-constrains the type, it defaults to the signed 32-bit integer i32.

r[expr.literal.int.inference-error]

• If the program context over-constrains the type, it is considered a static type error.

Examples of integer literal expressions:

123;	// type i32
123132;	// type i32
123u32;	// type u32
123_u32;	// type u32
let a: u64 = 123;	// type u64
0xff;	// type i32
0xff; 0xff_u8;	// type i32 // type u8
0xff; 0xff_u8;	// type i32 // type u8
0xff; 0xff_u8; 0o70;	// type i32 // type u8 // type i32
0xff; 0xff_u8; 0o70; 0o70_i16;	<pre>// type i32 // type u8 // type i32 // type i16</pre>

0b1111_1111_1001_0000;	// type i32
0b1111_1111_1001_0000i64;	// type i64
Ousize;	// type usize

r[expr.literal.int.representation] The value of the expression is determined from the string representation of the token as follows:

r[expr.literal.int.radix]

- An integer radix is chosen by inspecting the first two characters of the string, as follows:

  - 00 indicates radix 8
  - • indicates radix 16
  - otherwise the radix is 10.

r[expr.literal.int.radix-prefix-stripped]

• If the radix is not 10, the first two characters are removed from the string.

r[expr.literal.int.type-suffix-stripped]

• Any suffix is removed from the string. r[expr.literal.int.separators-stripped]

• Any underscores are removed from the string. r[expr.literal.int.u128-value]

The string is converted to a u128 value as if by
 [u128::from\_str\_radix] with the chosen radix. If the value does not
 fit in u128, it is a compiler error.

r[expr.literal.int.cast]

• The u128 value is converted to the expression's type via a <u>numeric</u> <u>cast</u>.

[!NOTE] The final cast will truncate the value of the literal if it does not fit in the expression's type. rustc includes a <u>lint check</u> named <u>overflowing\_literals</u>, defaulting to <u>deny</u>, which rejects expressions where this occurs.

[!NOTE] -118, for example, is an application of the <u>negation</u> <u>operator</u> to the literal expression 118, not a single integer literal expression. See <u>Overflow</u> for notes on representing the most negative value for a signed type.

r[expr.literal.float]

## **Floating-point literal expressions**

r[expr.literal.float.intro] A floating-point literal expression has one of two forms:

- a single [FLOAT\_LITERAL] token
- a single [INTEGER\_LITERAL] token which has a suffix and no radix indicator

r[expr.literal.float.suffix] If the token has a <u>suffix</u>, the suffix must be the name of one of the <u>primitive floating-point types</u>: f32 or f64, and the expression has that type.

r[expr.literal.float.infer] If the token has no suffix, the expression's type is determined by type inference:

r[expr.literal.float.inference-unique-type]

- If a floating-point type can be *uniquely* determined from the surrounding program context, the expression has that type. r[expr.literal.float.inference-default]
- If the program context under-constrains the type, it defaults to f64. r[expr.literal.float.inference-error]
- If the program context over-constrains the type, it is considered a static type error.

Examples of floating-point literal expressions:

123.0f64;	//	type	f64
0.1f64;	//	type	f64
0.1f32;	//	type	f32
12E+99_f64;	//	type	f64
5f32;	//	type	f32
let x: $f64 = 2.;$	//	type	f64

r[expr.literal.float.result] The value of the expression is determined from the string representation of the token as follows:

r[expr.literal.float.type-suffix-stripped]

- Any suffix is removed from the string. r[expr.literal.float.separators-stripped]
- Any underscores are removed from the string. r[expr.literal.float.value]
- The string is converted to the expression's type as if by <u>f32::from str</u> or <u>f64::from str</u>.

[!NOTE] -1.0, for example, is an application of the <u>negation</u> <u>operator</u> to the literal expression 1.0, not a single floating-point literal expression.

and literal tokens. The [!NOTE] inf NaN are not [f32::INFINITY], [f64::INFINITY], [f32::NAN], and [f64::NAN] constants can be used instead of literal expressions. In rustc, a literal large enough to be evaluated as infinite will trigger the overflowing\_literals lint check.

r[expr.literal.bool]

#### **Boolean literal expressions**

r[expr.literal.bool.intro] A boolean literal expression consists of one of the keywords true or false.

r[expr.literal.bool.result] The expression's type is the primitive <u>boolean</u> <u>type</u>, and its value is:

- true if the keyword is true
- false if the keyword is false

r[expr.path]

# **Path expressions**

r[expr.path.syntax]

PathExpression ->

PathInExpression

| QualifiedPathInExpression

r[expr.path.intro] A <u>path</u> used as an expression context denotes either a local variable or an item.

r[expr.path.place] Path expressions that resolve to local or static variables are <u>place expressions</u>, other paths are <u>value expressions</u>.

r[expr.path.safety] Using a static mut variable requires an unsafe
block.

```
# mod globals {
#     pub static STATIC_VAR: i32 = 5;
#     pub static mut STATIC_MUT_VAR: i32 = 7;
# }
# let local_var = 3;
local_var;
globals::STATIC_VAR;
unsafe { globals::STATIC_MUT_VAR };
let some_constructor = Some::<i32>;
let push_integer = Vec::<i32>::push;
let slice_reverse = <[i32]>::reverse;
```

r[expr.path.const] Evaluation of associated constants is handled the same way as <u>const\_blocks</u>.

r[expr.block]

# **Block expressions**

```
r[expr.block.syntax]
BlockExpression ->
    `{`
        InnerAttribute*
        Statements?
        `}`
```

Statements ->

Statement+

- | Statement+ ExpressionWithoutBlock
- | ExpressionWithoutBlock

r[expr.block.intro] A *block expression*, or *block*, is a control flow expression and anonymous namespace scope for items and variable declarations.

r[expr.block.sequential-evaluation] As a control flow expression, a block sequentially executes its component non-item declaration statements and then its final optional expression.

r[expr.block.namespace] As an anonymous namespace scope, item declarations are only in scope inside the block itself and variables declared by let statements are in scope from the next statement until the end of the block. See the <u>scopes</u> chapter for more details.

r[expr.block.inner-attributes] The syntax for a block is {, then any <u>inner</u> <u>attributes</u>, then any number of <u>statements</u>, then an optional expression, called the final operand, and finally a }.

r[expr.block.statements] Statements are usually required to be followed by a semicolon, with two exceptions:

- 1. Item declaration statements do not need to be followed by a semicolon.
- 2. Expression statements usually require a following semicolon except if its outer expression is a flow control expression.

r[expr.block.null-statement] Furthermore, extra semicolons between statements are allowed, but these semicolons do not affect semantics.

r[expr.block.evaluation] When evaluating a block expression, each statement, except for item declaration statements, is executed sequentially.

r[expr.block.result] Then the final operand is executed, if given.

r[expr.block.type] The type of a block is the type of the final operand, or() if the final operand is omitted.

```
# fn fn_call() {}
let _: () = {
    fn_call();
};
let five: i32 = {
    fn_call();
    5
};
assert_eq!(5, five);
```

[!NOTE] As a control flow expression, if a block expression is the outer expression of an expression statement, the expected type is () unless it is followed immediately by a semicolon.

r[expr.block.value] Blocks are always <u>value expressions</u> and evaluate the last operand in value expression context.

[!NOTE] This can be used to force moving a value if really needed. For example, the following example fails on the call to consume\_self
because the struct was moved out of s in the block expression.
struct Struct;

```
impl Struct {
    fn consume_self(self) {}
    fn borrow_self(&self) {}
}
```

```
fn move_by_block_expression() {
   let s = Struct;
   // Move the value out of `s` in the block expression.
   (&{ s }).borrow_self();
   // Fails to execute because `s` is moved out of.
   s.consume_self();
}
```

r[expr.block.async]
## async blocks

r[expr.block.async.syntax]
AsyncBlockExpression -> `async` `move`? BlockExpression

r[expr.block.async.intro] An *async block* is a variant of a block expression which evaluates to a future.

r[expr.block.async.future-result] The final expression of the block, if present, determines the result value of the future.

r[expr.block.async.anonymous-type] Executing an async block is similar to executing a closure expression: its immediate effect is to produce and return an anonymous type.

r[expr.block.async.future] Whereas closures return a type that implements one or more of the [std::ops::Fn] traits, however, the type returned for an async block implements the [std::future::Future] trait.

r[expr.block.async.layout-unspecified] The actual data format for this type is unspecified.

[!NOTE] The future type that rustc generates is roughly equivalent to an enum with one variant per await point, where each variant stores the data needed to resume from its corresponding point.

```
r[expr.block.async.edition2018]
```

[!EDITION-2018] Async blocks are only available beginning with Rust 2018.

r[expr.block.async.capture]

## **Capture modes**

Async blocks capture variables from their environment using the same <u>capture modes</u> as closures. Like closures, when written  $async \{ ... \}$  the capture mode for each variable will be inferred from the content of the block. async move  $\{ ... \}$  blocks however will move all referenced variables into the resulting future.

r[expr.block.async.context]

# Async context

Because async blocks construct a future, they define an **async context** which can in turn contain <u>await expressions</u>. Async contexts are established by async blocks as well as the bodies of async functions, whose semantics are defined in terms of async blocks.

r[expr.block.async.function]

## **Control-flow operators**

r[expr.block.async.function.intro] Async blocks act like a function boundary, much like closures.

r[expr.block.async.function.return-try] Therefore, the ? operator and return expressions both affect the output of the future, not the enclosing function or other context. That is, return <expr> from within an async block will return the result of <expr> as the output of the future. Similarly, if <expr>? propagates an error, that error is propagated as the result of the future.

r[expr.block.async.function.control-flow] Finally, the break and continue keywords cannot be used to branch out from an async block. Therefore the following is illegal:

```
loop {
    async move {
        break; // error[E0267]: `break` inside of an `async`
    block
     }
}
r[expr.block.const]
```

## const blocks

```
r[expr.block.const.syntax]
ConstBlockExpression -> `const` BlockExpression
```

r[expr.block.const.intro] A *const block* is a variant of a block expression whose body evaluates at compile-time instead of at runtime.

r[expr.block.const.context] Const blocks allows you to define a constant value without having to define new <u>constant items</u>, and thus they are also sometimes referred as *inline consts*. It also supports type inference so there is no need to specify the type, unlike <u>constant items</u>.

r[expr.block.const.generic-params] Const blocks have the ability to reference generic parameters in scope, unlike <u>free</u> constant items. They are desugared to constant items with generic parameters in scope (similar to associated constants, but without a trait or type they are associated with). For example, this code:

```
fn foo<T>() -> usize {
    const { std::mem::size_of::<T>() + 1 }
}
```

is equivalent to:

```
fn foo<T>() -> usize {
    {
        struct Const<T>(T);
        impl<T> Const<T> {
            const CONST: usize = std::mem::size_of::<T>() + 1;
        }
        Const::<T>::CONST
    }
}
```

```
r[expr.block.const.evaluation]
```

If the const block expression is executed at runtime, then the constant is guaranteed to be evaluated, even if its return value is ignored:

```
fn foo<T>() -> usize {
    // If this code ever gets executed, then the assertion has
```

```
definitely
   // been evaluated at compile-time.
   const { assert!(std::mem::size_of::<T>() > 0); }
   // Here we can have unsafe code relying on the type being
non-zero-sized.
   /* ... */
   42
}
```

r[expr.block.const.not-executed]

If the const block expression is not executed at runtime, it may or may not be evaluated:

```
if false {
```

// The panic may or may not occur when the program is built.

```
const { panic!(); }
```

}

r[expr.block.unsafe]

## unsafe blocks

```
r[expr.block.unsafe.syntax]
UnsafeBlockExpression -> `unsafe` BlockExpression
```

r[expr.block.unsafe.intro] *See* <u>unsafe</u> <u>blocks</u> for more information on when to use unsafe.

A block of code can be prefixed with the unsafe keyword to permit <u>unsafe operations</u>. Examples:

```
unsafe {
    let b = [13u8, 17u8];
    let a = &b[0] as *const u8;
    assert_eq!(*a, 13);
    assert_eq!(*a.offset(1), 17);
}
# unsafe fn an_unsafe_fn() -> i32 { 10 }
let a = unsafe { an_unsafe_fn() };
```

```
r[expr.block.label]
```

# Labelled block expressions

Labelled block expressions are documented in the <u>Loops and other</u> <u>breakable expressions</u> section.

r[expr.block.attributes]

# **Attributes on block expressions**

r[expr.block.attributes.inner-attributes] <u>Inner attributes</u> are allowed directly after the opening brace of a block expression in the following situations:

- <u>Function</u> and <u>method</u> bodies.
- Loop bodies (<u>loop</u>, <u>while</u>, and <u>for</u>).
- Block expressions used as a <u>statement</u>.
- Block expressions as elements of <u>array expressions</u>, <u>tuple expressions</u>, <u>call expressions</u>, and tuple-style <u>struct</u> expressions.
- A block expression as the tail expression of another block expression.

r[expr.block.attributes.valid] The attributes that have meaning on a block expression are <u>cfg</u> and <u>the lint check attributes</u>.

For example, this function returns true on unix platforms and false on other platforms.

```
fn is_unix_platform() -> bool {
    #[cfg(unix)] { true }
    #[cfg(not(unix))] { false }
}
```

r[expr.operator]

# **Operator expressions**

r[expr.operator.syntax]

OperatorExpression ->

BorrowExpression

| DereferenceExpression

| ErrorPropagationExpression

| NegationExpression

| ArithmeticOrLogicalExpression

| ComparisonExpression

| LazyBooleanExpression

| TypeCastExpression

| AssignmentExpression

| CompoundAssignmentExpression

r[expr.operator.intro] Operators are defined for built in types by the Rust language.

r[expr.operator.trait] Many of the following operators can also be overloaded using traits in std::ops or std::cmp.

r[expr.operator.int-overflow]

## Overflow

r[expr.operator.int-overflow.intro] Integer operators will panic when they overflow when compiled in debug mode. The -C debug-assertions and -C overflow-checks compiler flags can be used to control this more directly. The following things are considered to be overflow:

r[expr.operator.int-overflow.binary-arith]

• When +, \* or binary - create a value greater than the maximum value, or less than the minimum value that can be stored.

r[expr.operator.int-overflow.unary-neg]

• Applying unary - to the most negative value of any signed integer type, unless the operand is a <u>literal expression</u> (or a literal expression standing alone inside one or more <u>grouped</u> <u>expressions</u>).

r[expr.operator.int-overflow.div]

 Using / or %, where the left-hand argument is the smallest integer of a signed integer type and the right-hand argument is -1. These checks occur even when -C overflow-checks is disabled, for legacy reasons.

r[expr.operator.int-overflow.shift]

• Using << or >> where the right-hand argument is greater than or equal to the number of bits in the type of the left-hand argument, or is negative.

[!NOTE] The exception for literal expressions behind unary - means that forms such as -128\_i8 or let j: i8 = -(128) never cause a panic and have the expected value of -128.

In these cases, the literal expression already has the most negative value for its type (for example, <u>128\_i8</u> has the value -128) because integer literals are truncated to their type per the description in <u>Integer literal expressions</u>.

Negation of these most negative values leaves the value unchanged due to two's complement overflow conventions.

In rustc, these most negative expressions are also ignored by the overflowing\_literals lint check.

r[expr.operator.borrow]

#### **Borrow operators**

r[expr.operator.borrow.syntax]

```
BorrowExpression ->
(`&`|`&&`) Expression
| (`&`|`&&`) `mut` Expression
| (`&`|`&&`) `raw` `const` Expression
| (`&`|`&&`) `raw` `mut` Expression
```

r[expr.operator.borrow.intro] The & (shared borrow) and &mut (mutable borrow) operators are unary prefix operators.

r[expr.operator.borrow.result] When applied to a <u>place expression</u>, this expressions produces a reference (pointer) to the location that the value refers to.

r[expr.operator.borrow.lifetime] The memory location is also placed into a borrowed state for the duration of the reference. For a shared borrow (&), this implies that the place may not be mutated, but it may be read or shared again. For a mutable borrow (&mut), the place may not be accessed in any way until the borrow expires.

r[expr.operator.borrow.mut] &mut evaluates its operand in a mutable place expression context.

r[expr.operator.borrow.temporary] If the & or &mut operators are applied to a <u>value expression</u>, then a <u>temporary value</u> is created.

These operators cannot be overloaded.

```
{
    // a temporary with value 7 is created that lasts for this scope.
    let shared_reference = &7;
}
let mut array = [-2, 3, 9];
{
    // Mutably borrows `array` for this scope.
    // `array` may only be used through `mutable_reference`.
    let mutable_reference = &mut array;
}
```

r[expr.borrow.and-and-syntax] Even though && is a single token (<u>the lazy 'and' operator</u>), when used in the context of borrow expressions it works as two borrows:

```
// same meanings:
let a = && 10;
let a = & & 10;
// same meanings:
let a = &&&& mut 10;
let a = && && mut 10;
let a = & & & mut 10;
```

r[expr.borrow.raw]

### **Raw borrow operators**

r[expr.borrow.raw.intro] & raw const and & raw mut are the raw borrow operators.

r[expr.borrow.raw.place] The operand expression of these operators is evaluated in place expression context.

r[expr.borrow.raw.result] & raw const expr then creates a const raw pointer of type \* const T to the given place, and & raw mut expr creates a mutable raw pointer of type \* mut T.

r[expr.borrow.raw.invalid-ref] The raw borrow operators must be used instead of a borrow operator whenever the place expression could evaluate to a place that is not properly aligned or does not store a valid value as determined by its type, or whenever creating a reference would introduce incorrect aliasing assumptions. In those situations, using a borrow operator would cause <u>undefined</u> <u>behavior</u> by creating an invalid reference, but a raw pointer may still be constructed.

The following is an example of creating a raw pointer to an unaligned place through a packed struct:

```
#[repr(packed)]
struct Packed {
   f1: u8,
   f2: u16,
}
let packed = Packed { f1: 1, f2: 2 };
// `&packed.f2` would create an unaligned reference, and thus be undefined
behavior!
let raw_f2 = &raw const packed.f2;
assert_eq!(unsafe { raw_f2.read_unaligned() }, 2);
```

The following is an example of creating a raw pointer to a place that does not contain a valid value:

```
use std::mem::MaybeUninit;
struct Demo {
    field: bool,
}
let mut uninit = MaybeUninit::<Demo>::uninit();
// `&uninit.as_mut().field` would create a reference to an uninitialized `bool`,
// and thus be undefined behavior!
let f1_ptr = unsafe { &raw mut (*uninit.as_mut_ptr()).field };
unsafe { f1_ptr.write(true); }
let init = unsafe { uninit.assume_init() };
```

r[expr.deref]

#### The dereference operator

r[expr.deref.syntax]

DereferenceExpression -> `\*` Expression

r[expr.deref.intro] The \* (dereference) operator is also a unary prefix operator.

r[expr.deref.result] When applied to a pointer it denotes the pointed-to location.

r[expr.deref.mut] If the expression is of type &mut T or \*mut T, and is either a local variable, a (nested) field of a local variable or is a mutable <u>place expression</u>, then the resulting memory location can be assigned to.

r[expr.deref.safety] Dereferencing a raw pointer requires unsafe.

r[expr.deref.traits] On non-pointer types \*x is equivalent to \*std::ops::Deref::deref(&x) in an <u>immutable place expression context</u> and \*std::ops::DerefMut::deref\_mut(&mut x) in a mutable place expression context.

```
let x = &7;
assert_eq!(*x, 7);
let y = &mut 9;
*y = 11;
assert_eq!(*y, 11);
```

r[expr.try]

#### The question mark operator

r[expr.try.syntax]

ErrorPropagationExpression -> Expression `?`

r[expr.try.intro] The question mark operator (?) unwraps valid values or returns erroneous values, propagating them to the calling function.

r[expr.try.restricted-types] It is a unary postfix operator that can only be applied to the types Result<T, E> and Option<T>.

r[expr.try.behavior-std-result] When applied to values of the Result<T, E> type, it propagates errors.

r[expr.try.effects-err] If the value is Err(e), then it will return Err(From::from(e)) from the enclosing function or closure.

r[expr.try.result-ok] If applied to Ok(x), then it will unwrap the value to evaluate to x.

```
# use std::num::ParseIntError;
fn try_to_parse() -> Result<i32, ParseIntError> {
    let x: i32 = "123".parse()?; // x = 123
    let y: i32 = "24a".parse()?; // returns an Err() immediately
    Ok(x + y) // Doesn't run.
}
let res = try_to_parse();
println!("{:?}", res);
# assert!(res.is_err())
```

r[expr.try.behavior-std-option] When applied to values of the Option<T> type, it propagates None s.

r[expr.try.effects-none] If the value is None, then it will return None.

r[expr.try.result-some] If applied to Some(x), then it will unwrap the value to evaluate to x.

```
fn try_option_some() -> Option<u8> {
    let val = Some(1)?;
    Some(val)
}
assert_eq!(try_option_some(), Some(1));
fn try_option_none() -> Option<u8> {
    let val = None?;
    Some(val)
}
assert_eq!(try_option_none(), None);
r[expr.try.trait] ? cannot be overloaded.
r[expr.negate]
```

## **Negation operators**

r[expr.negate.syntax]

NegationExpression ->

`-` Expression

| `!` Expression

r[expr.negate.intro] These are the last two unary operators.

r[expr.negate.results] This table summarizes the behavior of them on primitive types and which traits are used to overload these operators for other types. Remember that signed integers are always represented using two's complement. The operands of all of these operators are evaluated in <u>value</u> expression context so are moved or copied.

Symbol	Integer	bool	<b>Floating Point</b>	<b>Overloading</b> Trait
-	Negation*		Negation	<pre>std::ops::Neg</pre>
!	Bitwise NOT	Logical NOT		<pre>std::ops::Not</pre>

\* Only for signed integer types.

Here are some example of these operators

```
let x = 6;
assert_eq!(-x, -6);
assert_eq!(!x, -7);
assert_eq!(true, !false);
```

r[expr.arith-logic]

## **Arithmetic and Logical Binary Operators**

r[expr.arith-logic.syntax]

```
ArithmeticOrLogicalExpression ->
```

	Expression	`+`	Expression
I	Expression	`-`	Expression
I	Expression	`*`	Expression
I	Expression	$^{\prime}/^{\prime}$	Expression
I	Expression	`%`	Expression
I	Expression	`&`	Expression
I	Expression	` `	Expression
I	Expression	`^`	Expression
I	Expression	`<<	` Expression
Τ	Expression	`>>	Expression

r[expr.arith-logic.intro] Binary operators expressions are all written with infix notation.

r[expr.arith-logic.behavior] This table summarizes the behavior of arithmetic and logical binary operators on primitive types and which traits are used to overload these operators for other types. Remember that signed integers are always represented using two's complement. The operands of all of these operators are evaluated in <u>value expression context</u> so are moved or copied.

Symbol	Integer	bool	Floating Point	Overloading Trait	Overloading Compound Assignment Trait
+	Addition		Addition	std::ops:: Add	std::ops:: AddAssign
•	Subtraction		Subtraction	std::ops:: Sub	std::ops:: SubAssign
*	Multiplication		Multiplication	std::ops:: Mul	std::ops:: MulAssign
/	Division*†		Division	std::ops:: Div	std::ops:: DivAssign
%	Remainder**†		Remainder	std::ops:: Rem	std::ops:: RemAssign
&	Bitwise AND	<u>Logical</u> <u>AND</u>		std::ops:: BitAnd	std::ops:: BitAndAssi gn
Ι	Bitwise OR	<u>Logical</u> <u>OR</u>		std::ops:: BitOr	std::ops:: BitOrAssig n

Symbol	Integer	bool	Floating Point	Overloading Trait	Overloading Compound Assignment Trait
^	Bitwise XOR	<u>Logical</u> <u>XOR</u>		std::ops:: BitXor	std::ops:: BitXorAssi gn
<<	Left Shift			std::ops:: Shl	std::ops:: ShlAssign
>>	Right Shift***			std::ops:: Shr	std::ops:: ShrAssign

\* Integer division rounds towards zero.

\*\* Rust uses a remainder defined with <u>truncating division</u>. Given remainder = dividend % divisor, the remainder will have the same sign as the dividend.

\*\*\* Arithmetic right shift on signed integer types, logical right shift on unsigned integer types.

† For integer types, division by zero panics.

Here are examples of these operators being used.

```
assert_eq!(3 + 6, 9);
assert_eq!(5.5 - 1.25, 4.25);
assert_eq!(-5 * 14, -70);
assert_eq!(14 / 3, 4);
assert_eq!(100 % 7, 2);
assert_eq!(0b1010 & 0b1100, 0b1000);
assert_eq!(0b1010 | 0b1100, 0b1110);
assert_eq!(0b1010 ^ 0b1100, 0b110);
assert_eq!(13 << 3, 104);
assert_eq!(-10 >> 2, -3);
```

r[expr.cmp]

### **Comparison Operators**

r[expr.cmp.syntax]

```
ComparisonExpression ->
```

```
Expression `==` Expression
| Expression `!=` Expression
| Expression `>` Expression
| Expression `<` Expression
| Expression `>=` Expression
| Expression `<=` Expression
```

r[expr.cmp.intro] Comparison operators are also defined both for primitive types and many types in the standard library.

r[expr.cmp.paren-chaining] Parentheses are required when chaining comparison operators. For example, the expression a == b == c is invalid and may be written as (a == b) == c.

r[expr.cmp.trait] Unlike arithmetic and logical operators, the traits for overloading these operators are used more generally to show how a type may be compared and will likely be assumed to define actual comparisons by functions that use these traits as bounds. Many functions and macros in the standard library can then use that assumption (although not to ensure safety).

r[expr.cmp.place] Unlike the arithmetic and logical operators above, these operators implicitly take shared borrows of their operands, evaluating them in <u>place expression context</u>:

```
# let a = 1;
# let b = 1;
a == b;
// is equivalent to
::std::cmp::PartialEq::eq(&a, &b);
```

This means that the operands don't have to be moved out of. r[expr.cmp.behavior]

Symbol	Meaning	Overloading method
==	Equal	<pre>std::cmp::PartialEq::eq</pre>
!=	Not equal	<pre>std::cmp::PartialEq::ne</pre>
>	Greater than	<pre>std::cmp::PartialOrd::gt</pre>
<	Less than	<pre>std::cmp::PartialOrd::lt</pre>
>=	Greater than or equal to	<pre>std::cmp::PartialOrd::ge</pre>
<=	Less than or equal to	<pre>std::cmp::PartialOrd::le</pre>

Here are examples of the comparison operators being used.

```
assert!(123 == 123);
assert!(23 != -12);
assert!(12.5 > 12.2);
assert!([1, 2, 3] < [1, 3, 4]);
assert!('A' <= 'B');
assert!("World" >= "Hello");
```

r[expr.bool-logic]

## Lazy boolean operators

r[expr.bool-logic.syntax]
LazyBooleanExpression ->
 Expression `||` Expression
 | Expression `&&` Expression

r[expr.bool-logic.intro] The operators || and && may be applied to operands of boolean type. The || operator denotes logical 'or', and the && operator denotes logical 'and'.

r[expr.bool-logic.conditional-evaluation] They differ from | and & in that the right-hand operand is only evaluated when the left-hand operand does not already determine the result of the expression. That is, || only evaluates its right-hand operand when the left-hand operand evaluates to false, and && only when it evaluates to true.

```
let x = false || true; // true
let y = false && panic!(); // false, doesn't evaluate `panic!()`
```

r[expr.as]

### **Type cast expressions**

r[expr.as.syntax]

TypeCastExpression -> Expression `as` TypeNoBounds

r[expr.as.intro] A type cast expression is denoted with the binary operator as.

r[expr.as.result] Executing an as expression casts the value on the left-hand side to the type on the right-hand side.

An example of an as expression:

```
# fn sum(values: &[f64]) -> f64 { 0.0 }
# fn len(values: &[f64]) -> i32 { 0 }
fn average(values: &[f64]) -> f64 {
    let sum: f64 = sum(values);
    let size: f64 = len(values) as f64;
    sum / size
}
```

r[expr.as.coercions] as can be used to explicitly perform <u>coercions</u>, as well as the following additional casts. Any cast that does not fit either a coercion rule or an entry in the table is a compiler error. Here **\*T** means either **\*const T** or **\*mut T**. **m** stands for optional **mut** in reference types and **mut** or const in pointer types.

Type of e	U	Cast performed by e as U	
Integer or Float type	Integer or Float type	[Numeric cast][expr.as.numeric]	
Enumeration	Integer type	[Enum cast][expr.as.enum]	
bool or char	Integer type	[Primitive to integer cast][expr.as.bool-char- as-int]	
u8	char	[u8 to char cast][expr.as.u8-as-char]	
*T	*V 1	[Pointer to pointer cast][expr.as.pointer]	
*⊤ where ⊤: Sized	Integer type	[Pointer to address cast][expr.as.pointer-as- int]	
Integer type	*V where V: Sized	[Address to pointer cast][expr.as.int-as- pointer]	
&m1 [T; n]	*m <sub>2</sub> T <sup>2</sup>	Array to pointer cast	
*m1 [T; n]	*m <sub>2</sub> T <sup>2</sup>	Array to pointer cast	
Function item	Function pointer	Function item to function pointer cast	
Function item	*V where V: Sized	Function item to pointer cast	
Function item	Integer	Function item to address cast	

Type of e	U	Cast performed by e as U
Function pointer	*V where V: Sized	Function pointer to pointer cast
Function pointer	Integer	Function pointer to address cast
Closure <sup>3</sup>	Function pointer	Closure to function pointer cast

```
1
```

where  $\top$  and  $\lor$  have compatible metadata:

- V: Sized, or
- Both slice metadata (\*[u16] -> \*[u8], \*str -> \*(u8, [u32])), or
- Both the same trait object metadata, modulo dropping auto traits (\*dyn Debug -> \*(u16, dyn Debug), \*dyn Debug + Send -> \*dyn Debug)
  - Note: adding auto traits is only allowed if the principal trait has the auto trait as a super trait (given trait T: Send {}, \*dyn T -> \*dyn T + Send is valid, but \*dyn Debug -> \*dyn Debug + Send is not)
  - Note: Generics (including lifetimes) must match (\*dyn T<'a, A> -> \*dyn T<'b, B> requires 'a = 'b and A = B)

2

3

only when  $m_1$  is mut or  $m_2$  is const. Casting mut reference/pointer to const pointer is allowed.

only for closures that do not capture (close over) any local variables can be casted to function pointers.

## Semantics

r[expr.as.numeric]

#### Numeric cast

r[expr.as.numeric.int-same-size]

• Casting between two integers of the same size (e.g. i32 -> u32) is a no-op (Rust uses 2's complement for negative values of fixed integers)

```
assert_eq!(42i8 as u8, 42u8);
assert_eq!(-1i8 as u8, 255u8);
assert_eq!(255u8 as i8, -1i8);
assert_eq!(-1i16 as u16, 65535u16);
```

r[expr.as.numeric.int-truncation]

 Casting from a larger integer to a smaller integer (e.g. u32 -> u8) will truncate assert\_eq!(42u16 as u8, 42u8); assert\_eq!(1234u16 as u8, 210u8); assert\_eq!(0xabcdu16 as u8, 0xcdu8);

```
assert_eq!(-42i16 as i8, -42i8);
assert_eq!(1234u16 as i8, -46i8);
assert_eq!(0xabcdi32 as i8, -51i8);
```

r[expr.as.numeric.int-extension]

• Casting from a smaller integer to a larger integer (e.g. u8 -> u32) will

```
o zero-extend if the source is unsigned
o sign-extend if the source is signed
assert_eq!(42i8 as i16, 42i16);
assert_eq!(-17i8 as i16, -17i16);
assert_eq!(0b1000_1010u8 as u16, 0b0000_0000_1000_1010u16, "Zero-extend");
assert_eq!(0b0000_1010i8 as i16, 0b0000_0000_0000_1010i16, "Sign-extend 0");
assert_eq!(0b1000_1010u8 as i8 as i16, 0b1111_1111_1000_1010u16 as i16, "Sign-
extend 1");
```

r[expr.as.numeric.float-as-int]

- Casting from a float to an integer will round the float towards zero
  - NaN will return 0
  - Values larger than the maximum integer value, including **INFINITY**, will saturate to the maximum value of the integer type.
  - Values smaller than the minimum integer value, including **NEG\_INFINITY**, will saturate to the minimum value of the integer type.

```
assert_eq!(42.9f32 as i32, 42);
assert_eq!(-42.9f32 as i32, -42);
assert_eq!(42_000_000f32 as i32, 42_000_000);
assert_eq!(std::f32::NAN as i32, 0);
assert_eq!(1_000_000_000_000_000f32 as i32, 0x7fffffffi32);
assert_eq!(std::f32::NEG_INFINITY as i32, -0x80000000i32);
```

r[expr.as.numeric.int-as-float]

- Casting from an integer to float will produce the closest possible float \*
  - if necessary, rounding is according to roundTiesToEven mode \*\*\*
  - on overflow, infinity (of the same sign as the input) is produced
  - note: with the current set of numeric types, overflow can only happen on u128 as f32 for values greater or equal to f32::MAX + (0.5 ULP)

r[expr.as.numeric.float-widening]

• Casting from an f32 to an f64 is perfect and lossless

```
assert_eq!(1_234.5f32 as f64, 1_234.5f64);
assert_eq!(std::f32::INFINITY as f64, std::f64::INFINITY);
assert!((std::f32::NAN as f64).is_nan());
```

r[expr.as.numeric.float-narrowing]

Casting from an f64 to an f32 will produce the closest possible f32 \*\*

```
    if necessary, rounding is according to roundTiesToEven mode ***
    on overflow, infinity (of the same sign as the input) is produced
    assert_eq!(1_234.5f64 as f32, 1_234.5f32);
    assert_eq!(1_234_567_891.123f64 as f32, 1_234_567_890f32, "Rounded");
    assert_eq!(std::f64::INFINITY as f32, std::f32::INFINITY);
    assert!((std::f64::NAN as f32).is_nan());
```

\* if integer-to-float casts with this rounding mode and overflow behavior are not supported natively by the hardware, these casts will likely be slower than expected.

\*\* if f64-to-f32 casts with this rounding mode and overflow behavior are not supported natively by the hardware, these casts will likely be slower than expected.

\*\*\* as defined in IEEE 754-2008 §4.3.1: pick the nearest floating point number, preferring the one with an even least significant digit if exactly halfway between two floating point numbers.

r[expr.as.enum]

#### Enum cast

r[expr.as.enum.discriminant] Casts an enum to its discriminant, then uses a numeric cast if needed. Casting is limited to the following kinds of enumerations:

```
    <u>Unit-only enums</u>
```

• <u>Field-less enums</u> without <u>explicit discriminants</u>, or where only unit-variants have explicit discriminants

```
enum Enum { A, B, C }
assert_eq!(Enum::A as i32, 0);
assert_eq!(Enum::B as i32, 1);
assert_eq!(Enum::C as i32, 2);
```

r[expr.as.enum.no-drop] Casting is not allowed if the enum implements [Drop].

r[expr.as.bool-char-as-int]

#### **Primitive to integer cast**

- false casts to 0, true casts to 1
- char casts to the value of the code point, then uses a numeric cast if needed.

```
assert_eq!(false as i32, 0);
assert_eq!(true as i32, 1);
assert_eq!('A' as i32, 65);
assert_eq!('Ö' as i32, 214);
```

r[expr.as.u8-as-char]

u8 to char cast

Casts to the char with the corresponding code point.

assert\_eq!(65u8 as char, 'A'); assert\_eq!(214u8 as char, 'O');

r[expr.as.pointer-as-int]

#### Pointer to address cast

Casting from a raw pointer to an integer produces the machine address of the referenced memory. If the integer type is smaller than the pointer type, the address may be truncated; using usize avoids this.

r[expr.as.int-as-pointer]

#### Address to pointer cast

Casting from an integer to a raw pointer interprets the integer as a memory address and produces a pointer referencing that memory.

[!WARNING] This interacts with the Rust memory model, which is still under development. A pointer obtained from this cast may suffer additional restrictions even if it is bitwise equal to a valid pointer. Dereferencing such a pointer may be <u>undefined behavior</u> if aliasing rules are not followed.

A trivial example of sound address arithmetic:

```
let mut values: [i32; 2] = [1, 2];
let p1: *mut i32 = values.as_mut_ptr();
let first_address = p1 as usize;
let second_address = first_address + 4; // 4 == size_of::<i32>()
let p2 = second_address as *mut i32;
unsafe {
    *p2 += 1;
}
assert_eq!(values[1], 3);
```

r[expr.as.pointer]

#### **Pointer-to-pointer cast**

r[expr.as.pointer.behavior] \*const T / \*mut T can be cast to \*const U / \*mut U with the following behavior:

r[expr.as.pointer.sized]

• If  $\top$  and  $\cup$  are both sized, the pointer is returned unchanged.

r[expr.as.pointer.unsized]

• If T and U are both unsized, the pointer is also returned unchanged. In particular, the metadata is preserved exactly.

For instance, a cast from \*const [T] to \*const [U] preserves the number of elements. Note that, as a consequence, such casts do not necessarily preserve the size of the pointer's referent (e.g., casting \*const [u16] to \*const [u8] will result in a raw pointer which refers to an object of half the size of the original). The same holds for str and any compound type whose unsized tail is a slice type, such as struct Foo(i32, [u8]) or (u64, Foo). r[expr.as.pointer.discard-metadata]

If T is unsized and U is sized, the cast discards all metadata that completes the wide pointer T and produces a thin pointer U consisting of the data part of the unsized pointer.
 r[expr.assign]

#### **Assignment expressions**

r[expr.assign.syntax]

AssignmentExpression -> Expression `=` Expression

r[expr.assign.intro] An assignment expression moves a value into a specified place.

r[expr.assign.assignee] An assignment expression consists of a <u>mutable assignee expression</u>, the *assignee operand*, followed by an equals sign (=) and a <u>value expression</u>, the *assigned value operand*.

r[expr.assign.behavior-basic] In its most basic form, an assignee expression is a <u>place expression</u>, and we discuss this case first.

r[expr.assign.behavior-destructuring] The more general case of destructuring assignment is discussed below, but this case always decomposes into sequential assignments to place expressions, which may be considered the more fundamental case.

r[expr.assign.basic]

#### **Basic assignments**

r[expr.assign.evaluation-order] Evaluating assignment expressions begins by evaluating its operands. The assigned value operand is evaluated first, followed by the assignee expression.

r[expr.assign.destructuring-order] For destructuring assignment, subexpressions of the assignee expression are evaluated left-to-right.

[!NOTE] This is different than other expressions in that the right operand is evaluated before the left one.

r[expr.assign.drop-target] It then has the effect of first <u>dropping</u> the value at the assigned place, unless the place is an uninitialized local variable or an uninitialized field of a local variable.

r[expr.assign.behavior] Next it either <u>copies or moves</u> the assigned value to the assigned place.

r[expr.assign.result] An assignment expression always produces the unit value.

Example:

let mut x = 0; let y = 0; x = y;

r[expr.assign.destructure]

#### **Destructuring assignments**

r[expr.assign.destructure.intro] Destructuring assignment is a counterpart to destructuring pattern matches for variable declaration, permitting assignment to complex values, such as tuples or structs. For instance, we may swap two mutable variables:

let (mut a, mut b) = (0, 1);
// Swap `a` and `b` using destructuring assignment.
(b, a) = (a, b);

r[expr.assign.destructure.assignee] In contrast to destructuring declarations using let, patterns may not appear on the left-hand side of an assignment due to syntactic ambiguities. Instead, a group of expressions that correspond to patterns are designated to be <u>assignee expressions</u>, and permitted on

the left-hand side of an assignment. Assignee expressions are then desugared to pattern matches followed by sequential assignment.

r[expr.assign.destructure.irrefutable] The desugared patterns must be irrefutable: in particular, this means that only slice patterns whose length is known at compile-time, and the trivial slice [...], are permitted for destructuring assignment.

The desugaring method is straightforward, and is illustrated best by example.

```
# struct Struct { x: u32, y: u32 }
# let (mut a, mut b) = (0, 0);
(a, b) = (3, 4);
[a, b] = [3, 4];
Struct { x: a, y: b } = Struct { x: 3, y: 4};
// desugars to:
{
    let (_a, _b) = (3, 4);
    a = _a;
    b = b;
}
{
    let [_a, _b] = [3, 4];
    a = _a;
    b = b;
}
{
    let Struct { x: _a, y: _b } = Struct { x: 3, y: 4};
    a = _a;
    b = b;
}
```

r[expr.assign.destructure.repeat-ident] Identifiers are not forbidden from being used multiple times in a single assignee expression.

r[expr.assign.destructure.discard-value] <u>Underscore expressions</u> and empty <u>range expressions</u> may be used to ignore certain values, without binding them.

r[expr.assign.destructure.default-binding] Note that default binding modes do not apply for the desugared expression.

r[expr.compound-assign]

#### **Compound assignment expressions**

```
r[expr.compound-assign.syntax]
```

```
CompoundAssignmentExpression ->
Expression `+=` Expression
| Expression `-=` Expression
| Expression `*=` Expression
| Expression `/=` Expression
| Expression `&=` Expression
| Expression `A=` Expression
| Expression `^=` Expression
| Expression `<=` Expression
| Expression `<=` Expression
```

r[expr.compound-assign.intro] *Compound assignment expressions* combine arithmetic and logical binary operators with assignment expressions.

For example:

```
let mut x = 5;
x += 1;
assert!(x == 6);
```

The syntax of compound assignment is a <u>mutable place expression</u>, the *assigned operand*, then one of the operators followed by an = as a single token (no whitespace), and then a <u>value expression</u>, the *modifying operand*.

r[expr.compound-assign.place] Unlike other place operands, the assigned place operand must be a place expression.

r[expr.compound-assign.no-value] Attempting to use a value expression is a compiler error rather than promoting it to a temporary.

r[expr.compound-assign.operand-order] Evaluation of compound assignment expressions depends on the types of the operators.

r[expr.compound-assign.primitive-order] If both types are primitives, then the modifying operand will be evaluated first followed by the assigned operand. It will then set the value of the assigned operand's place to the value of performing the operation of the operator with the values of the assigned operand and modifying operand.

[!NOTE] This is different than other expressions in that the right operand is evaluated before the left one.

r[expr.compound-assign.trait] Otherwise, this expression is syntactic sugar for calling the function of the overloading compound assignment trait of the operator (see the table earlier in this chapter). A mutable borrow of the assigned operand is automatically taken.

For example, the following expression statements in example are equivalent:

```
# struct Addable;
# use std::ops::AddAssign;
impl AddAssign<Addable> for Addable {
    /* */
```

```
# fn add_assign(&mut self, other: Addable) {}
}
fn example() {
# let (mut a1, a2) = (Addable, Addable);
a1 += a2;
# let (mut a1, a2) = (Addable, Addable);
AddAssign::add_assign(&mut a1, a2);
}
```

r[expr.compound-assign.result] Like assignment expressions, compound assignment expressions always produce <u>the unit value</u>.

[!WARNING] The evaluation order of operands swaps depending on the types of the operands: with primitive types the right-hand side will get evaluated first, while with non-primitive types the left-hand side will get evaluated first. Try not to write code that depends on the evaluation order of operands in compound assignment expressions. See <u>this test</u> for an example of using this dependency.

r[expr.paren]

# **Grouped expressions**

r[expr.paren.syntax]

GroupedExpression -> `(` Expression `)`

r[expr.paren.intro] A *parenthesized expression* wraps a single expression, evaluating to that expression. The syntax for a parenthesized expression is a ( , then an expression, called the *enclosed operand*, and then a ).

r[expr.paren.evaluation] Parenthesized expressions evaluate to the value of the enclosed operand.

r[expr.paren.place-or-value] Unlike other expressions, parenthesized expressions are both <u>place expressions and value expressions</u>. When the enclosed operand is a place expression, it is a place expression and when the enclosed operand is a value expression, it is a value expression.

r[expr.paren.override-precedence] Parentheses can be used to explicitly modify the precedence order of subexpressions within an expression.

An example of a parenthesized expression:

```
let x: i32 = 2 + 3 * 4; // not parenthesized
let y: i32 = (2 + 3) * 4; // parenthesized
assert_eq!(x, 14);
assert_eq!(y, 20);
```

An example of a necessary use of parentheses is when calling a function pointer that is a member of a struct:

```
# struct A {
# f: fn() -> &'static str
# }
# impl A {
# fn f(&self) -> &'static str {
# The method f"
# }
# }
# }
# let a = A{f: || "The field f"};
#
```

assert\_eq!( a.f (), "The method f"); assert\_eq!((a.f)(), "The field f"); r[expr.array]

# Array and array index expressions

## **Array expressions**

```
r[expr.array.syntax]
ArrayExpression -> `[` ArrayElements? `]`
```

```
ArrayElements ->
```

```
Expression ( `,` Expression )* `,`?
| Expression `;` Expression
```

r[expr.array.constructor] *Array expressions* construct <u>arrays</u>. Array expressions come in two forms.

r[expr.array.array] The first form lists out every value in the array.

r[expr.array.array-syntax] The syntax for this form is a comma-separated list of expressions of uniform type enclosed in square brackets.

r[expr.array.array-behavior] This produces an array containing each of these values in the order they are written.

r[expr.array.repeat] The syntax for the second form is two expressions separated by a semicolon (;) enclosed in square brackets.

r[expr.array.repeat-operand] The expression before the ; is called the *repeat operand*.

r[expr.array.length-operand] The expression after the ; is called the *length operand*.

r[expr.array.length-restriction] The length operand must either be an <u>inferred const</u> or be a <u>constant expression</u> of type usize (e.g. a <u>literal</u> or a <u>constant item</u>).

```
const C: usize = 1;
let _: [u8; C] = [0; 1]; // Literal.
let _: [u8; C] = [0; C]; // Constant item.
let _: [u8; C] = [0; _]; // Inferred const.
let _: [u8; C] = [0; (((_)))]; // Inferred const.
```

[!NOTE] In an array expression, an <u>inferred const</u> is parsed as an [expression][Expression] but then semantically treated as a separate kind of <u>const generic argument</u>.
r[expr.array.repeat-behavior] An array expression of this form creates an array with the length of the value of the length operand with each element being a copy of the repeat operand. That is, [a; b] creates an array containing b copies of the value of a.

r[expr.array.repeat-copy] If the length operand has a value greater than 1 then this requires that the type of the repeat operand is <u>Copy</u> or that it must be a <u>path</u> to a constant item.

r[expr.array.repeat-const-item] When the repeat operand is a constant item, it is evaluated the length operand's value times.

r[expr.array.repeat-evaluation-zero] If that value is <sup>0</sup>, then the constant item is not evaluated at all.

r[expr.array.repeat-non-const] For expressions that are not a constant item, it is evaluated exactly once, and then the result is copied the length operand's value times.

## Array and slice indexing expressions

r[expr.array.index.syntax]
IndexExpression -> Expression `[` Expression `]`

r[expr.array.index.array] <u>Array</u> and <u>slice</u>-typed values can be indexed by writing a square-bracket-enclosed expression of type <u>usize</u> (the index) after them. When the array is mutable, the resulting <u>memory location</u> can be assigned to.

r[expr.array.index.trait] For other types an index expression a[b] is equivalent to \*std::ops::Index::index(&a, b), or \*std::ops::IndexMut::index\_mut(&mut a, b) in a mutable place expression context. Just as with methods, Rust will also insert dereference operations on a repeatedly to find an implementation.

r[expr.array.index.zero-index] Indices are zero-based for arrays and slices.

r[expr.array.index.const] Array access is a <u>constant expression</u>, so bounds can be checked at compile-time with a constant index value. Otherwise a check will be performed at run-time that will put the thread in a <u>panicked state</u> if it fails.

```
// lint is deny by default.
#![warn(unconditional_panic)]
```

```
([1, 2, 3, 4])[2]; // Evaluates to 3
let b = [[1, 0, 0], [0, 1, 0], [0, 0, 1]];
b[1][2]; // multidimensional array indexing
let x = (["a", "b"])[10]; // warning: index out of bounds
let n = 10;
let y = (["a", "b"])[n]; // panics
let arr = ["a", "b"];
arr[10]; // warning: index out of bounds
```

r[expr.array.index.trait-impl] The array index expression can be implemented for types other than arrays and slices by implementing the <u>Index</u> and <u>IndexMut</u> traits.

r[expr.tuple]

## Tuple and tuple indexing expressions

#### **Tuple expressions**

r[expr.tuple.syntax]

TupleExpression -> `(` TupleElements? `)`

TupleElements -> ( Expression `,` )+ Expression?

r[expr.tuple.result] A *tuple expression* constructs <u>tuple values</u>.

r[expr.tuple.intro] The syntax for tuple expressions is a parenthesized, comma separated list of expressions, called the *tuple initializer operands*.

r[expr.tuple.unary-tuple-restriction] 1-ary tuple expressions require a comma after their tuple initializer operand to be disambiguated with a <u>parenthetical expression</u>.

r[expr.tuple.value] Tuple expressions are a <u>value expression</u> that evaluate into a newly constructed value of a tuple type.

r[expr.tuple.type] The number of tuple initializer operands is the arity of the constructed tuple.

r[expr.tuple.unit] Tuple expressions without any tuple initializer operands produce the unit tuple.

r[expr.tuple.fields] For other tuple expressions, the first written tuple initializer operand initializes the field 0 and subsequent operands initializes the next highest field. For example, in the tuple expression ('a', 'b', 'c'), 'a' initializes the value of the field 0, 'b' field 1, and 'c' field 2.

Examples of tuple expressions and their types:

Expression	Туре
()	() (unit)
(0.0, 4.5)	(f64, f64)
<pre>("x".to_string(), )</pre>	(String, )
("a", 4usize, true)	(&'static str, usize, bool)

r[expr.tuple-index]

## **Tuple indexing expressions**

r[expr.tuple-index.syntax]

TupleIndexingExpression -> Expression `.` TUPLE\_INDEX

r[expr.tuple-index.intro] A *tuple indexing expression* accesses fields of <u>tuples</u> and <u>tuple structs</u>.

The syntax for a tuple index expression is an expression, called the *tuple operand*, then a  $\cdot$ , then finally a tuple index.

r[expr.tuple-index.index-syntax] The syntax for the *tuple index* is a <u>decimal literal</u> with no leading zeros, underscores, or suffix. For example 0 and 2 are valid tuple indices but not 01, 0\_, nor 0i32.

r[expr.tuple-index.required-type] The type of the tuple operand must be a <u>tuple type</u> or a <u>tuple struct</u>.

r[expr.tuple-index.index-name-operand] The tuple index must be a name of a field of the type of the tuple operand.

r[expr.tuple-index.result] Evaluation of tuple index expressions has no side effects beyond evaluation of its tuple operand. As a <u>place expression</u>, it evaluates to the location of the field of the tuple operand with the same name as the tuple index.

Examples of tuple indexing expressions:

```
// Indexing a tuple
let pair = ("a string", 2);
assert_eq!(pair.1, 2);
// Indexing a tuple struct
# struct Point(f32, f32);
let point = Point(1.0, 0.0);
assert_eq!(point.0, 1.0);
assert_eq!(point.1, 0.0);
```

[!NOTE] Unlike field access expressions, tuple index expressions can be the function operand of a <u>call expression</u> as it cannot be confused with a method call since method names cannot be numbers.

[!NOTE] Although arrays and slices also have elements, you must use an <u>array or slice indexing expression</u> or a <u>slice pattern</u> to access their elements. r[expr.struct]

## **Struct expressions**

```
r[expr.struct.syntax]
StructExpression ->
    PathInExpression `{` (StructExprFields | StructBase)? `}`
StructExprFields ->
    StructExprField (`,` StructExprField)* (`,` StructBase |
`,`?)
StructExprField ->
    OuterAttribute*
    (
        IDENTIFIER
        | (IDENTIFIER | TUPLE_INDEX) `:` Expression
    )
```

```
StructBase -> `..` Expression
```

r[expr.struct.intro] A *struct expression* creates a struct, enum, or union value. It consists of a path to a <u>struct</u>, <u>enum variant</u>, or <u>union</u> item followed by the values for the fields of the item.

The following are examples of struct expressions:

```
# struct Point { x: f64, y: f64 }
# struct NothingInMe { }
# mod game { pub struct User<'a> { pub name: &'a str, pub age:
u32, pub score: usize } }
# enum Enum { Variant {} }
Point {x: 10.0, y: 20.0};
NothingInMe {};
let u = game::User {name: "Joe", age: 35, score: 100_000};
Enum::Variant {};
```

[!NOTE] Tuple structs and tuple enum variants are typically instantiated using a [call expression][expr.call] referring to the

[constructor in the value namespace][items.struct.tuple]. These are distinct from a struct expression using curly braces referring to the constructor in the type namespace.

```
struct Position(i32, i32, i32);
Position(0, 0, 0); // Typical way of creating a tuple
struct.
let c = Position; // `c` is a function that takes 3
arguments.
let pos = c(8, 6, 7); // Creates a `Position` value.
enum Version { Triple(i32, i32, i32) };
Version::Triple(0, 0, 0);
let f = Version::Triple;
let ver = f(8, 6, 7);
```

The last segment of the call path cannot refer to a type alias:

```
trait Tr { type T; }
impl<T> Tr for T { type T = T; }
struct Tuple();
enum Enum { Tuple() }
// <Unit as Tr>::T(); // causes an error -- `::T` is a
type, not a value
<Enum as Tr>::T::Tuple(); // OK
```

Unit structs and unit enum variants are typically instantiated using a [path expression][expr.path] referring to the [constant in the value namespace][items.struct.unit].

```
struct Gamma;
// Gamma unit value, referring to the const in the value
namespace.
let a = Gamma;
// Exact same value as `a`, but constructed using a
struct expression
```

```
// referring to the type namespace.
let b = Gamma {};
enum ColorSpace { Oklch }
let c = ColorSpace::Oklch;
let d = ColorSpace::Oklch {};
```

r[expr.struct.field]

## **Field struct expression**

r[expr.struct.field.intro] A struct expression with fields enclosed in curly braces allows you to specify the value for each individual field in any order. The field name is separated from its value with a colon.

r[expr.struct.field.union-constraint] A value of a <u>union</u> type can only be created using this syntax, and it must specify exactly one field.

r[expr.struct.update]

### **Functional update syntax**

r[expr.struct.update.intro] A struct expression that constructs a value of a struct type can terminate with the syntax . followed by an expression to denote a functional update.

r[expr.struct.update.base-same-type] The expression following . . (the base) must have the same struct type as the new struct type being formed.

r[expr.struct.update.fields] The entire expression uses the given values for the fields that were specified and moves or copies the remaining fields from the base expression.

r[expr.struct.update.visibility-constraint] As with all struct expressions, all of the fields of the struct must be <u>visible</u>, even those not explicitly named.

```
# struct Point3d { x: i32, y: i32, z: i32 }
let mut base = Point3d {x: 1, y: 2, z: 3};
let y_ref = &mut base.y;
Point3d {y: 0, z: 10, .. base}; // OK, only base.x is accessed
drop(y_ref);
```

r[expr.struct.brace-restricted-positions] Struct expressions can't be used directly in a <u>loop</u> or <u>if</u> expression's head, or in the <u>scrutinee</u> of an <u>if let</u> or <u>match</u> expression. However, struct expressions can be used in these situations if they are within another expression, for example inside <u>parentheses</u>.

r[expr.struct.tuple-field] The field names can be decimal integer values to specify indices for constructing tuple structs. This can be used with base structs to fill out the remaining indices not specified:

struct Color(u8, u8, u8); let c1 = Color(0, 0, 0); // Typical way of creating a tuple struct. let c2 = Color{0: 255, 1: 127, 2: 0}; // Specifying fields by index. let c3 = Color{1: 0, ..c2}; // Fill out all other fields using a base struct. r[expr.struct.field.named]

## Struct field init shorthand

When initializing a data structure (struct, enum, union) with named (but not numbered) fields, it is allowed to write fieldname as a shorthand for fieldname: fieldname. This allows a compact syntax with less duplication. For example:

```
# struct Point3d { x: i32, y: i32, z: i32 }
# let x = 0;
# let y_value = 0;
# let z = 0;
Point3d { x: x, y: y_value, z: z };
Point3d { x, y: y_value, z };
```

r[expr.call]

## **Call expressions**

```
r[expr.call.syntax]
CallExpression -> Expression `(` CallParams? `)`
```

```
CallParams -> Expression ( `,` Expression )* `,`?
```

r[expr.call.intro] A *call expression* calls a function. The syntax of a call expression is an expression, called the *function operand*, followed by a parenthesized comma-separated list of expression, called the *argument operands*.

r[expr.call.convergence] If the function eventually returns, then the expression completes.

r[expr.call.trait] For <u>non-function types</u>, the expression f(...) uses the method on one of the following traits based on the function operand:

- [Fn] or [AsyncFn] --- shared reference.
- [FnMut] or [AsyncFnMut] --- mutable reference.
- [FnOnce] or [AsyncFnOnce] --- value.

r[expr.call.autoref-deref] An automatic borrow will be taken if needed. The function operand will also be <u>automatically dereferenced</u> as required.

Some examples of call expressions:

```
# fn add(x: i32, y: i32) -> i32 { 0 }
let three: i32 = add(1i32, 2i32);
let name: &'static str = (|| "Rust")();
```

r[expr.call.desugar]

## **Disambiguating Function Calls**

r[expr.call.desugar.fully-qualified] All function calls are sugar for a more explicit <u>fully-qualified syntax</u>.

r[expr.call.desugar.ambiguity] Function calls may need to be fully qualified, depending on the ambiguity of a call in light of in-scope items.

[!NOTE] In the past, the terms "Unambiguous Function Call Syntax", "Universal Function Call Syntax", or "UFCS", have been used in documentation, issues, RFCs, and other community writings. However, these terms lack descriptive power and potentially confuse the issue at hand. We mention them here for searchability's sake.

r[expr.call.desugar.limits] Several situations often occur which result in ambiguities about the receiver or referent of method or associated function calls. These situations may include:

- Multiple in-scope traits define methods with the same name for the same types
- Auto-deref is undesirable; for example, distinguishing between methods on a smart pointer itself and the pointer's referent
- Methods which take no arguments, like <u>default()</u>, and return properties of a type, like <u>size of()</u>

r[expr.call.desugar.explicit-path] To resolve the ambiguity, the programmer may refer to their desired method or function using more specific paths, types, or traits.

For example,

```
trait Pretty {
    fn print(&self);
}
trait Ugly {
    fn print(&self);
}
```

```
struct Foo;
impl Pretty for Foo {
    fn print(&self) {}
}
struct Bar;
impl Pretty for Bar {
    fn print(&self) {}
}
impl Ugly for Bar {
    fn print(&self) {}
}
fn main() {
    let f = Foo;
    let b = Bar;
     // we can do this because we only have one item called
`print` for `Foo`s
    f.print();
   // more explicit, and, in the case of `Foo`, not necessary
    Foo::print(&f);
    // if you're not into the whole brevity thing
    <Foo as Pretty>::print(&f);
    // b.print(); // Error: multiple 'print' found
     // Bar::print(&b); // Still an error: multiple `print`
found
    // necessary because of in-scope items defining `print`
    <Bar as Pretty>::print(&b);
```

```
}
```

Refer to <u>RFC 132</u> for further details and motivations.

r[expr.method]

## **Method-call expressions**

r[expr.method.syntax]

MethodCallExpression -> Expression `.` PathExprSegment
`(`CallParams? `)`

r[expr.method.intro] A *method call* consists of an expression (the *receiver*) followed by a single dot, an expression path segment, and a parenthesized expression-list.

r[expr.method.target] Method calls are resolved to associated <u>methods</u> on specific traits, either statically dispatching to a method if the exact self-type of the left-hand-side is known, or dynamically dispatching if the left-hand-side expression is an indirect <u>trait object</u>.

```
let pi: Result<f32, _> = "3.14".parse();
let log_pi = pi.unwrap_or(1.0).log(2.72);
# assert!(1.14 < log_pi && log_pi < 1.15)</pre>
```

r[expr.method.autoref-deref] When looking up a method call, the receiver may be automatically dereferenced or borrowed in order to call a method. This requires a more complex lookup process than for other functions, since there may be a number of possible methods to call. The following procedure is used:

r[expr.method.candidate-receivers] The first step is to build a list of candidate receiver types. Obtain these by repeatedly <u>dereferencing</u> the receiver expression's type, adding each type encountered to the list, then finally attempting an <u>unsized coercion</u> at the end, and adding the result type if that is successful.

r[expr.method.candidate-receivers-refs] Then, for each candidate  $\top$ , add &T and &mut T to the list immediately after T.

For instance, if the receiver has type Box<[i32;2]>, then the candidate types will be Box<[i32;2]>, &Box<[i32;2]>, &mut Box<[i32;2]>, [i32; 2] (by dereferencing), &[i32; 2], &mut [i32; 2], [i32] (by unsized coercion), &[i32], and finally &mut [i32].

r[expr.method.candidate-search] Then, for each candidate type T, search for a <u>visible</u> method with a receiver of that type in the following places:

- 1. T's inherent methods (methods implemented directly on T).
- 2. Any of the methods provided by a <u>visible</u> trait implemented by τ. If τ is a type parameter, methods provided by trait bounds on τ are looked up first. Then all remaining methods in scope are looked up.

[!NOTE] The lookup is done for each type in order, which can occasionally lead to surprising results. The below code will print "In trait impl!", because <code>&self</code> methods are looked up first, the trait method is found before the struct's <code>&mut self</code> method is found.

```
struct Foo {}
trait Bar {
  fn bar(&self);
}
impl Foo {
  fn bar(&mut self) {
    println!("In struct impl!")
  }
}
impl Bar for Foo {
  fn bar(&self) {
    println!("In trait impl!")
  }
}
fn main() {
  let mut f = Foo{};
  f.bar();
}
```

r[expr.method.ambiguous-target] If this results in multiple possible candidates, then it is an error, and the receiver must be <u>converted</u> to an appropriate receiver type to make the method call.

r[expr.method.receiver-constraints] This process does not take into account the mutability or lifetime of the receiver, or whether a method is unsafe. Once a method is looked up, if it can't be called for one (or more) of those reasons, the result is a compiler error.

r[expr.method.ambiguous-search] If a step is reached where there is more than one possible method, such as where generic methods or traits are considered the same, then it is a compiler error. These cases require a <u>disambiguating function call syntax</u> for method and function invocation.

r[expr.method.edition2021]

[!EDITION-2021] Before the 2021 edition, during the search for visible methods, if the candidate receiver type is an <u>array type</u>, methods provided by the standard library <u>IntoIterator</u> trait are ignored.

The edition used for this purpose is determined by the token representing the method name.

This special case may be removed in the future.

[!WARNING] For <u>trait objects</u>, if there is an inherent method of the same name as a trait method, it will give a compiler error when trying to call the method in a method call expression. Instead, you can call the method using <u>disambiguating function call syntax</u>, in which case it calls the trait method, not the inherent method. There is no way to call the inherent method. Just don't define inherent methods on trait objects with the same name as a trait method and you'll be fine.

r[expr.field]

## **Field access expressions**

r[expr.field.syntax]

FieldExpression -> Expression `.` IDENTIFIER

r[expr.field.intro] A *field expression* is a <u>place expression</u> that evaluates to the location of a field of a <u>struct</u> or <u>union</u>.

r[expr.field.mut] When the operand is <u>mutable</u>, the field expression is also mutable.

r[expr.field.form] The syntax for a field expression is an expression, called the *container operand*, then a \_, and finally an <u>identifier</u>.

r[expr.field.not-method-call] Field expressions cannot be followed by a parenthetical comma-separated list of expressions, as that is instead parsed as a <u>method call expression</u>. That is, they cannot be the function operand of a <u>call expression</u>.

[!NOTE] Wrap the field expression in a <u>parenthesized expression</u> to use it in a call expression.

```
# struct HoldsCallable<F: Fn()> { callable: F }
let holds_callable = HoldsCallable { callable: || () };
```

```
// Invalid: Parsed as calling the method "callable"
// holds_callable.callable();
```

```
// Valid
(holds_callable.callable)();
```

```
Examples:
mystruct.myfield;
foo().x;
(Struct {a: 10, b: 20}).a;
(mystruct.function_field)() // Call expression containing a
field expression
```

```
r[expr.field.autoref-deref]
```

#### **Automatic dereferencing**

If the type of the container operand implements **Deref** or **DerefMut** depending on whether the operand is **mutable**, it is *automatically dereferenced* as many times as necessary to make the field access possible. This process is also called *autoderef* for short.

r[expr.field.borrow]

#### Borrowing

The fields of a struct or a reference to a struct are treated as separate entities when borrowing. If the struct does not implement <u>Drop</u> and is stored in a local variable, this also applies to moving out of each of its fields. This also does not apply if automatic dereferencing is done through user-defined types other than <u>Box</u>.

```
struct A { f1: String, f2: String, f3: String }
let mut x: A;
\# x = A \{
      f1: "f1".to_string(),
#
      f2: "f2".to_string(),
#
      f3: "f3".to_string()
#
# };
let a: &mut String = &mut x.f1; // x.f1 borrowed mutably
let b: \&String = \&x.f2;
                            // x.f2 borrowed immutably
let c: &String = &x.f2;
                               // Can borrow again
let d: String = x.f3;
                                // Move out of x.f3
```

r[expr.closure]

## **Closure expressions**

```
r[expr.closure.syntax]
ClosureExpression ->
`async`?[^cl-async-edition]
`move`?
(`||` | `|` ClosureParameters? `|` )
(Expression | `->` TypeNoBounds BlockExpression)
ClosureParameters -> ClosureParam (`,` ClosureParam)* `,`?
ClosureParam -> OuterAttribute* PatternNoTopAlt (`:` Type )?
1
```

The async qualifier is not allowed in the 2015 edition.

r[expr.closure.intro] A *closure expression*, also known as a lambda expression or a lambda, defines a <u>closure type</u> and evaluates to a value of that type. The syntax for a closure expression is an optional <u>async</u> keyword, an optional <u>move</u> keyword, then a pipe-symbol-delimited (|) comma-separated list of <u>patterns</u>, called the *closure parameters* each optionally followed by a : and a type, then an optional -> and type, called the *return type*, and then an expression, called the *closure body operand*.

r[expr.closure.param-type] The optional type after each pattern is a type annotation for the pattern.

r[expr.closure.explicit-type-body] If there is a return type, the closure body must be a <u>block</u>.

r[expr.closure.parameter-restriction] A closure expression denotes a function that maps a list of parameters onto the expression that follows the parameters. Just like a <u>let binding</u>, the closure parameters are irrefutable <u>patterns</u>, whose type annotation is optional and will be inferred from context if not given.

r[expr.closure.unique-type] Each closure expression has a unique, anonymous type.

r[expr.closure.captures] Significantly, closure expressions *capture their environment*, which regular <u>function definitions</u> do not.

r[expr.closure.capture-inference] Without the move keyword, the closure expression <u>infers how it captures each variable from its environment</u>, preferring to capture by shared reference, effectively borrowing all outer variables mentioned inside the closure's body.

r[expr.closure.capture-mut-ref] If needed the compiler will infer that instead mutable references should be taken, or that the values should be moved or copied (depending on their type) from the environment.

r[expr.closure.capture-move] A closure can be forced to capture its environment by copying or moving values by prefixing it with the move keyword. This is often used to ensure that the closure's lifetime is 'static'.

r[expr.closure.trait-impl]

#### **Closure trait implementations**

Which traits the closure type implement depends on how variables are captured, the types of the captured variables, and the presence of async. See the <u>call traits and coercions</u> chapter for how and when a closure implements Fn, FnMut, and FnOnce. The closure type implements <u>Send</u> and <u>Sync</u> if the type of every captured variable also implements the trait.

r[expr.closure.async]

## Async closures

r[expr.closure.async.intro] Closures marked with the async keyword indicate that they are asynchronous in an analogous way to an [async function][items.fn.async].

r[expr.closure.async.future] Calling the async closure does not perform any work, but instead evaluates to a value that implements [Future] that corresponds to the computation of the body of the closure.

```
async fn takes_async_callback(f: impl AsyncFn(u64)) {
   f(0).await;
   f(1).await;
}
async fn example() {
   takes_async_callback(async |i| {
      core::future::ready(i).await;
      println!("done with {i}.");
   }).await;
}
```

r[expr.closure.async.edition2018]

[!EDITION-2018] Async closures are only available beginning with Rust 2018.

### Example

In this example, we define a function ten\_times that takes a higherorder function argument, and we then call it with a closure expression as an argument, followed by a closure expression that moves values from its environment.

```
fn ten_times<F>(f: F) where F: Fn(i32) {
    for index in 0..10 {
        f(index);
    }
}
ten_times(|j| println!("hello, {}", j));
// With type annotations
ten_times(|j: i32| -> () { println!("hello, {}", j) });
let word = "konnichiwa".to_owned();
ten_times(move |j| println!("{}, {}", word, j));
```

## **Attributes on closure parameters**

r[expr.closure.param-attributes] Attributes on closure parameters follow the same rules and restrictions as <u>regular function parameters</u>.

r[expr.loop]

# Loops and other breakable expressions

```
r[expr.loop.syntax]
```

```
LoopExpression ->
```

```
LoopLabel? (
```

InfiniteLoopExpression

- | PredicateLoopExpression
- | IteratorLoopExpression
- | LabelBlockExpression

)

r[expr.loop.intro] Rust supports four loop expressions:

- A <u>loop expression</u> denotes an infinite loop.
- A <u>while expression</u> loops until a predicate is false.
- A <u>for expression</u> extracts values from an iterator, looping until the iterator is empty.
- A <u>labelled block expression</u> runs a loop exactly once, but allows exiting the loop early with break.

r[expr.loop.break-label] All four types of loop support <u>break</u> <u>expressions</u>, and <u>labels</u>.

r[expr.loop.continue-label] All except labelled block expressions support <u>continue expressions</u>.

r[expr.loop.explicit-result] Only loop and labelled block expressions support <u>evaluation to non-trivial values</u>.

r[expr.loop.infinite]
## **Infinite loops**

r[expr.loop.infinite.syntax]

InfiniteLoopExpression -> `loop` BlockExpression

r[expr.loop.infinite.intro] A loop expression repeats execution of its body continuously: loop { println!("I live."); }.

r[expr.loop.infinite.diverging] A loop expression without an associated break expression is diverging and has type <u>!</u>.

r[expr.loop.infinite.break] A loop expression containing associated <u>break expression(s)</u> may terminate, and must have type compatible with the value of the break expression(s).

r[expr.loop.while]

# **Predicate loops**

r[expr.loop.while.grammar]
PredicateLoopExpression -> `while` Conditions BlockExpression

r[expr.loop.while.intro] A while loop expression allows repeating the evaluation of a block while a set of conditions remain true.

r[expr.loop.while.syntax] The syntax of a while expression is a sequence of one or more condition operands separated by &&, followed by a [BlockExpression].

r[expr.loop.while.condition] Condition operands must be either an [Expression] with a <u>boolean type</u> or a conditional let match. If all of the condition operands evaluate to true and all of the let patterns successfully match their <u>scrutinees</u>, then the loop body block executes.

r[expr.loop.while.repeat] After the loop body successfully executes, the condition operands are re-evaluated to determine if the body should be executed again.

r[expr.loop.while.exit] If any condition operand evaluates to false or any let pattern does not match its scrutinee, the body is not executed and execution continues after the while expression.

r[expr.loop.while.eval] A while expression evaluates to ().

An example:

```
let mut i = 0;
while i < 10 {
    println!("hello");
    i = i + 1;
}
```

r[expr.loop.while.let]

### while let patterns

r[expr.loop.while.let.intro] let patterns in a while condition allow binding new variables into scope when the pattern matches successfully.

The following examples illustrate bindings using let patterns:

```
let mut x = vec![1, 2, 3];
while let Some(y) = x.pop() {
    println!("y = {}", y);
}
while let _ = 5 {
    println!("Irrefutable patterns are always true");
    break;
}
```

r[expr.loop.while.let.desugar] A while let loop is equivalent to a loop expression containing a <u>match expression</u> as follows.

```
'label: while let PATS = EXPR {
    /* loop body */
}
    is equivalent to
'label: loop {
    match EXPR {
        PATS => { /* loop body */ },
        _ => break,
    }
}
```

r[expr.loop.while.let.or-pattern] Multiple patterns may be specified with the | operator. This has the same semantics as with | in match expressions:

```
let mut vals = vec![2, 3, 1, 2, 2];
while let Some(v @ 1) | Some(v @ 2) = vals.pop() {
    // Prints 2, 2, then 1
    println!("{}", v);
}
```

r[expr.loop.while.chains]

### while condition chains

r[expr.loop.while.chains.intro] Multiple condition operands can be separated with &&. These have the same semantics and restrictions as <u>if</u> <u>condition chains</u>.

The following is an example of chaining multiple expressions, mixing **let** bindings and boolean expressions, and with expressions able to reference pattern bindings from previous expressions:

```
fn main() {
    let outer_opt = Some(Some(1i32));
    while let Some(inner_opt) = outer_opt
        && let Some(number) = inner_opt
        && number == 1
        {
            println!("Peek a boo");
            break;
        }
}
```

r[expr.loop.for]

## **Iterator loops**

```
r[expr.loop.for.syntax]
```

IteratorLoopExpression ->

`for` Pattern `in` Expression \_except [StructExpression]\_ BlockExpression

r[expr.loop.for.intro] A for expression is a syntactic construct for looping over elements provided by an implementation of std::iter::IntoIterator.

r[expr.loop.for.condition] If the iterator yields a value, that value is matched against the irrefutable pattern, the body of the loop is executed, and then control returns to the head of the for loop. If the iterator is empty, the for expression completes.

An example of a for loop over the contents of an array:

```
let v = &["apples", "cake", "coffee"];
for text in v {
    println!("I like {}.", text);
}
```

An example of a for loop over a series of integers:

```
let mut sum = 0;
for n in 1..11 {
    sum += n;
}
assert_eq!(sum, 55);
```

r[expr.loop.for.desugar] A for loop is equivalent to a loop expression containing a <u>match expression</u> as follows:

```
'label: for PATTERN in iter_expr {
    /* loop body */
}
```

is equivalent to

{
 let result = match IntoIterator::into\_iter(iter\_expr) {
 mut iter => 'label: loop {
 let mut next;
 match Iterator::next(&mut iter) {
 Option::Some(val) => next = val,
 Option::None => break,
 };
 let PATTERN = next;
 let () = { /\* loop body \*/ };
 },
 };
 result
}

r[expr.loop.for.lang-items] IntoIterator, Iterator, and Option are always the standard library items here, not whatever those names resolve to in the current scope.

The variable names next, iter, and val are for exposition only, they do not actually have names the user can type.

[!NOTE] The outer match is used to ensure that any <u>temporary</u> <u>values</u> in <u>iter\_expr</u> don't get dropped before the loop is finished. next is declared before being assigned because it results in types being inferred correctly more often.

r[expr.loop.label]

# **Loop labels**

```
r[expr.loop.label.syntax]
LoopLabel -> LIFETIME_OR_LABEL `:`
```

r[expr.loop.label.intro] A loop expression may optionally have a label. The label is written as a lifetime preceding the loop expression, as in 'foo: loop { break 'foo; }, 'bar: while false {}, 'humbug: for \_ in 0..0 {}.

r[expr.loop.label.control-flow] If a label is present, then labeled break and continue expressions nested within this loop may exit out of this loop or return control to its head. See <u>break expressions</u> and <u>continue</u> <u>expressions</u>.

r[expr.loop.label.ref] Labels follow the hygiene and shadowing rules of local variables. For example, this code will print "outer loop":

```
'a: loop {
    'a: loop {
        break 'a;
    }
    print!("outer loop");
    break 'a;
}
```

'\_ is not a valid loop label.

r[expr.loop.break]

### break expressions

```
r[expr.loop.break.syntax]
```

BreakExpression -> `break` LIFETIME\_OR\_LABEL? Expression?

r[expr.loop.break.intro] When break is encountered, execution of the associated loop body is immediately terminated, for example:

```
let mut last = 0;
for x in 1..100 {
    if x > 12 {
        break;
    }
    last = x;
}
assert_eq!(last, 12);
```

r[expr.loop.break.label] A break expression is normally associated with the innermost loop, for or while loop enclosing the break expression, but a <u>label</u> can be used to specify which enclosing loop is affected. Example:

```
'outer: loop {
    while true {
        break 'outer;
    }
}
```

r[expr.loop.break.value] A break expression is only permitted in the body of a loop, and has one of the forms break, break 'label or (see below) break EXPR or break 'label EXPR.

r[expr.loop.block-labels]

# Labelled block expressions

```
r[expr.loop.block-labels.syntax]
LabelBlockExpression -> BlockExpression
```

r[expr.loop.block-labels.intro] Labelled block expressions are exactly like block expressions, except that they allow using break expressions within the block.

r[expr.loop.block-labels.break] Unlike loops, break expressions within a labelled block expression *must* have a label (i.e. the label is not optional).

r[expr.loop.block-labels.label-required] Similarly, labelled block expressions *must* begin with a label.

```
# fn do_thing() {}
# fn condition_not_met() -> bool { true }
# fn do_next_thing() {}
# fn do_last_thing() {}
let result = 'block: {
    do_thing();
    if condition_not_met() {
        break 'block 1;
    }
    do_next_thing();
    if condition_not_met() {
        break 'block 2;
    }
    do_last_thing();
    3
};
 r[expr.loop.continue]
```

#### continue expressions

r[expr.loop.continue.syntax]
ContinueExpression -> `continue` LIFETIME\_OR\_LABEL?

r[expr.loop.continue.intro] When **continue** is encountered, the current iteration of the associated loop body is immediately terminated, returning control to the loop *head*.

r[expr.loop.continue.while] In the case of a while loop, the head is the conditional operands controlling the loop.

r[expr.loop.continue.for] In the case of a for loop, the head is the callexpression controlling the loop.

r[expr.loop.continue.label] Like break, continue is normally
associated with the innermost enclosing loop, but continue 'label may
be used to specify the loop affected.

r[expr.loop.continue.in-loop-only] A continue expression is only permitted in the body of a loop.

r[expr.loop.break-value]

### break and loop values

r[expr.loop.break-value.intro] When associated with a loop, a break expression may be used to return a value from that loop, via one of the forms break EXPR or break 'label EXPR, where EXPR is an expression whose result is returned from the loop. For example:

```
let (mut a, mut b) = (1, 1);
let result = loop {
    if b > 10 {
        break b;
    }
    let c = a + b;
    a = b;
    b = c;
};
// first number in Fibonacci sequence over 10:
assert_eq!(result, 13);
```

r[expr.loop.break-value.loop] In the case a loop has an associated break, it is not considered diverging, and the loop must have a type compatible with each break expression. break without an expression is considered identical to break with expression ().

r[expr.range]

# **Range expressions**

r[expr.range.syntax]
RangeExpression ->
 RangeExpr
 RangeFromExpr
 RangeToExpr
 RangeFullExpr
 RangeInclusiveExpr
 RangeToInclusiveExpr
RangeFromExpr -> Expression `..` Expression
RangeToExpr -> `..` Expression
RangeFullExpr -> `..`
RangeInclusiveExpr -> Expression `..=` Expression

RangeToInclusiveExpr -> `..=` Expression

r[expr.range.behavior] The .. and ..= operators will construct an object of one of the std::ops::Range (or core::ops::Range) variants, according to the following table:

Production	Syntax	Туре	Range
[RangeExpr]	start . . end	[std::ops::Range]	start ≤ x < end
[RangeFromExpr]	start	[std::ops::RangeFrom]	start ≤ x
[RangeToExpr]	end	[std::ops::RangeTo]	x < end
[RangeFullExpr]		[std::ops::RangeFull]	-
[RangeInclusiveExpr]	start . .= end	[std::ops::RangeInclusive]	$\begin{array}{l} \text{start} \leq \\ x \leq \\ \text{end} \end{array}$

Production	Syntax	Туре	Range
[RangeToInclusiveExpr]	= en d	[std::ops::RangeToInclusive]	x ≤ end

Examples:

12;	<pre>// std::ops::Range</pre>
3;	<pre>// std::ops::RangeFrom</pre>
4;	<pre>// std::ops::RangeTo</pre>
;	<pre>// std::ops::RangeFull</pre>
5=6;	<pre>// std::ops::RangeInclusive</pre>
=7;	<pre>// std::ops::RangeToInclusive</pre>

r[expr.range.equivalence] The following expressions are equivalent.

```
let x = std::ops::Range {start: 0, end: 10};
let y = 0..10;
assert_eq!(x, y);
r[expr.range.for] Ranges can be used in for loops:
for i in 1..11 {
    println!("{}", i);
}
```

r[expr.if]

# if expressions

r[expr.if.syntax] IfExpression -> `if` Conditions BlockExpression (`else` ( BlockExpression | IfExpression ) )? Conditions -> Expression \_except [StructExpression]\_ | LetChain LetChain -> LetChainCondition ( `&&` LetChainCondition )\* LetChainCondition -> Expression \_except [ExcludedConditions]\_ | OuterAttribute\* `let` Pattern `=` Scrutinee \_except [ExcludedConditions]\_ @root ExcludedConditions -> **StructExpression** | LazyBooleanExpression | RangeExpr | RangeFromExpr | RangeInclusiveExpr | AssignmentExpression | CompoundAssignmentExpression

r[expr.if.intro] The syntax of an if expression is a sequence of one or more condition operands separated by &&, followed by a consequent block, any number of else if conditions and blocks, and an optional trailing else block.

r[expr.if.condition] Condition operands must be either an [Expression] with a <u>boolean type</u> or a conditional let match.

r[expr.if.condition-true] If all of the condition operands evaluate to true and all of the let patterns successfully match their <u>scrutinees</u>, the consequent block is executed and any subsequent else if or else block is skipped.

r[expr.if.else-if] If any condition operand evaluates to false or any let pattern does not match its scrutinee, the consequent block is skipped and any subsequent else if condition is evaluated.

r[expr.if.else] If all if and else if conditions evaluate to false then any else block is executed.

r[expr.if.result] An if expression evaluates to the same value as the executed block, or () if no block is evaluated.

r[expr.if.type] An if expression must have the same type in all situations.

```
# let x = 3;
if x == 4 {
    println!("x is four");
} else if x == 3 {
    println!("x is three");
} else {
    println!("x is something else");
}
// `if` can be used as an expression.
let y = if 12 * 15 > 150 {
    "Bigger"
} else {
    "Smaller"
};
assert_eq!(y, "Bigger");
```

r[expr.if.let]

### if let patterns

r[expr.if.let.intro] let patterns in an if condition allow binding new variables into scope when the pattern matches successfully.

The following examples illustrate bindings using let patterns:

```
let dish = ("Ham", "Eggs");
// This body will be skipped because the pattern is refuted.
if let ("Bacon", b) = dish {
    println!("Bacon is served with {}", b);
} else {
    // This block is evaluated instead.
    println!("No bacon will be served");
}
// This body will execute.
if let ("Ham", b) = dish {
    println!("Ham is served with {}", b);
}
if let _ = 5 {
    println!("Irrefutable patterns are always true");
}
```

r[expr.if.let.or-pattern] Multiple patterns may be specified with the | operator. This has the same semantics as with | in <u>match expressions</u>:

```
enum E {
    X(u8),
    Y(u8),
    Z(u8),
}
let v = E::Y(12);
if let E::X(n) | E::Y(n) = v {
    assert_eq!(n, 12);
}
```

r[expr.if.chains]

# **Chains of conditions**

r[expr.if.chains.intro] Multiple condition operands can be separated with && .

r[expr.if.chains.order] Similar to a && [LazyBooleanExpression], each operand is evaluated from left-to-right until an operand evaluates as false or a let match fails, in which case the subsequent operands are not evaluated.

r[expr.if.chains.bindings] The bindings of each pattern are put into scope to be available for the next condition operand and the consequent block.

The following is an example of chaining multiple expressions, mixing let bindings and boolean expressions, and with expressions able to reference pattern bindings from previous expressions:

```
fn single() {
    let outer_opt = Some(Some(1i32));
    if let Some(inner_opt) = outer_opt
        && let Some(number) = inner_opt
        && number == 1
        {
            println!("Peek a boo");
        }
}
```

The above is equivalent to the following without using chains of conditions:

```
fn nested() {
    let outer_opt = Some(Some(1i32));
    if let Some(inner_opt) = outer_opt {
        if let Some(number) = inner_opt {
            if number == 1 {
                println!("Peek a boo");
            }
        }
    }
}
```

} }

r[expr.if.chains.or] If any condition operand is a let pattern, then none of the condition operands can be a [] [lazy boolean operator expression] [expr.bool-logic] due to ambiguity and precedence with the let scrutinee. If a [] expression is needed, then parentheses can be used. For example:

```
# let foo = Some(123);
# let condition1 = true;
# let condition2 = false;
// Parentheses are required here.
if let Some(x) = foo && (condition1 || condition2) { /*...*/ }
```

r[expr.if.edition2024]

[!EDITION-2024] Before the 2024 edition, let chains are not supported. That is, the [LetChain] grammar is not allowed in an if expression.

r[expr.match]

# match expressions

```
r[expr.match.syntax]
MatchExpression ->
  `match` Scrutinee `{`
    InnerAttribute*
    MatchArms?
  `}`
Scrutinee -> Expression _except [StructExpression]_
MatchArms ->
        ( MatchArm `=>` ( ExpressionWithoutBlock `,` |
ExpressionWithBlock `,`? ) )*
    MatchArm `=>` Expression `,`?
MatchArm -> OuterAttribute* Pattern MatchArmGuard?
```

MatchArmGuard -> `if` Expression

r[expr.match.intro] A *match expression* branches on a pattern. The exact form of matching that occurs depends on the <u>pattern</u>.

r[expr.match.scrutinee] A match expression has a *scrutinee* expression, which is the value to compare to the patterns.

r[expr.match.scrutinee-constraint] The scrutinee expression and the patterns must have the same type.

r[expr.match.scrutinee-behavior] A match behaves differently depending on whether or not the scrutinee expression is a <u>place expression</u> or <u>value expression</u>.

r[expr.match.scrutinee-value] If the scrutinee expression is a <u>value</u> <u>expression</u>, it is first evaluated into a temporary location, and the resulting value is sequentially compared to the patterns in the arms until a match is found. The first arm with a matching pattern is chosen as the branch target

of the match, any variables bound by the pattern are assigned to local variables in the arm's block, and control enters the block.

r[expr.match.scrutinee-place] When the scrutinee expression is a <u>place</u> <u>expression</u>, the match does not allocate a temporary location; however, a by-value binding may copy or move from the memory location. When possible, it is preferable to match on place expressions, as the lifetime of these matches inherits the lifetime of the place expression rather than being restricted to the inside of the match.

An example of a match expression:

```
let x = 1;
match x {
    1 => println!("one"),
    2 => println!("two"),
    3 => println!("two"),
    4 => println!("three"),
    4 => println!("four"),
    5 => println!("five"),
    _ => println!("something else"),
}
```

r[expr.match.pattern-vars] Variables bound within the pattern are scoped to the match guard and the arm's expression.

r[expr.match.pattern-var-binding] The <u>binding mode</u> (move, copy, or reference) depends on the pattern.

r[expr.match.or-pattern] Multiple match patterns may be joined with theoperator. Each pattern will be tested in left-to-right sequence until a successful match is found.

```
let x = 9;
let message = match x {
    0 | 1 => "not many",
    2 ..= 9 => "a few",
    _ => "lots"
};
assert_eq!(message, "a few");
```

```
// Demonstration of pattern match order.
struct S(i32, i32);
match S(1, 2) {
    S(z @ 1, _) | S(_, z @ 2) => assert_eq!(z, 1),
    _ => panic!(),
}
```

[!NOTE] The 2..=9 is a <u>Range Pattern</u>, not a <u>Range Expression</u>. Thus, only those types of ranges supported by range patterns can be used in match arms.

r[expr.match.or-patterns-restriction] Every binding in each | separated pattern must appear in all of the patterns in the arm.

r[expr.match.binding-restriction] Every binding of the same name must have the same type, and have the same binding mode.

r[expr.match.guard]

# Match guards

r[expr.match.guard.intro] Match arms can accept *match guards* to further refine the criteria for matching a case.

r[expr.match.guard.type] Pattern guards appear after the pattern and consist of a bool-typed expression following the if keyword.

r[expr.match.guard.behavior] When the pattern matches successfully, the pattern guard expression is executed. If the expression evaluates to true, the pattern is successfully matched against.

r[expr.match.guard.next] Otherwise, the next pattern, including other matches with the | operator in the same arm, is tested.

```
# let maybe_digit = Some(0);
# fn process_digit(i: i32) { }
# fn process_other(i: i32) { }
let message = match maybe_digit {
    Some(x) if x < 10 => process_digit(x),
    Some(x) => process_other(x),
    None => panic!(),
};
```

[!NOTE] Multiple matches using the poperator can cause the pattern guard and the side effects it has to execute multiple times. For example:

```
# use std::cell::Cell;
let i : Cell<i32> = Cell::new(0);
match 1 {
    1 | _ if { i.set(i.get() + 1); false } => {}
    _ => {}
}
assert_eq!(i.get(), 2);
```

r[expr.match.guard.bound-variables] A pattern guard may refer to the variables bound within the pattern they follow.

r[expr.match.guard.shared-ref] Before evaluating the guard, a shared reference is taken to the part of the scrutinee the variable matches on. While

evaluating the guard, this shared reference is then used when accessing the variable.

r[expr.match.guard.value] Only when the guard evaluates to true is the value moved, or copied, from the scrutinee into the variable. This allows shared borrows to be used inside guards without moving out of the scrutinee in case guard fails to match.

r[expr.match.guard.no-mutation] Moreover, by holding a shared reference while evaluating the guard, mutation inside guards is also prevented.

r[expr.match.attributes]

## **Attributes on match arms**

r[expr.match.attributes.outer] Outer attributes are allowed on match arms. The only attributes that have meaning on match arms are cfg and the <u>lint check attributes</u>.

r[expr.match.attributes.inner] <u>Inner attributes</u> are allowed directly after the opening brace of the match expression in the same expression contexts as <u>attributes on block expressions</u>. r[expr.return]

# return expressions

r[expr.return.syntax]

```
ReturnExpression -> `return` Expression?
```

r[expr.return.intro] Return expressions are denoted with the keyword return.

r[expr.return.behavior] Evaluating a return expression moves its argument into the designated output location for the current function call, destroys the current function activation frame, and transfers control to the caller frame.

An example of a return expression:

```
fn max(a: i32, b: i32) -> i32 {
    if a > b {
        return a;
    }
    return b;
}
```

r[expr.await]

# **Await expressions**

r[expr.await.syntax]

AwaitExpression -> Expression `.``await`

r[expr.await.intro] An await expression is a syntactic construct for suspending a computation provided by an implementation of std::future::IntoFuture until the given future is ready to produce a value.

r[expr.await.construct] The syntax for an await expression is an expression with a type that implements the IntoFuture trait, called the *future operand*, then the token ., and then the await keyword.

r[expr.await.allowed-positions] Await expressions are legal only within an <u>async context</u>, like an <u>async fn</u>, <u>async closure</u>, or <u>async block</u>.

r[expr.await.effects] More specifically, an await expression has the following effect.

- 1. Create a future by calling <u>IntoFuture::into future</u> on the future operand.
- 2. Evaluate the future to a <u>future</u> tmp;
- 3. Pin tmp using Pin::new unchecked;
- 4. This pinned future is then polled by calling the Future::poll method
   and passing it the current task context;
- 5. If the call to poll returns <u>Poll::Pending</u>, then the future returns Poll::Pending, suspending its state so that, when the surrounding async context is re-polled, execution returns to step 3;
- 6. Otherwise the call to poll must have returned <u>Poll::Ready</u>, in which case the value contained in the <u>Poll::Ready</u> variant is used as the result of the await expression itself.

r[expr.await.edition2018]

[!EDITION-2018] Await expressions are only available beginning with Rust 2018.

r[expr.await.task]

### **Task context**

The task context refers to the <u>Context</u> which was supplied to the current <u>async context</u> when the async context itself was polled. Because <u>await</u> expressions are only legal in an async context, there must be some task context available.

r[expr.await.desugar]

## **Approximate desugaring**

Effectively, an await expression is roughly equivalent to the following non-normative desugaring:

```
match operand.into_future() {
    mut pinned => loop {
        let mut pin = unsafe { Pin::new_unchecked(&mut pinned)
};
        match Pin::future::poll(Pin::borrow(&mut pin), &mut
current_context) {
        Poll::Ready(r) => break r,
        Poll::Pending => yield Poll::Pending,
        }
    }
}
```

where the yield pseudo-code returns Poll::Pending and, when reinvoked, resumes execution from that point. The variable current\_context refers to the context taken from the async environment. r[expr.placeholder]
# \_ expressions

r[expr.placeholder.syntax]
UnderscoreExpression -> `\_`

r[expr.placeholder.intro] Underscore expressions, denoted with the symbol \_\_, are used to signify a placeholder in a destructuring assignment.

r[expr.placeholder.lhs-assignment-only] They may only appear in the left-hand side of an assignment.

r[expr.placeholder.pattern] Note that this is distinct from the <u>wildcard</u> <u>pattern</u>.

```
Examples of _____expressions:
```

```
let p = (1, 2);
let mut a = 0;
(_, a) = p;
struct Position {
    x: u32,
    y: u32,
}
Position { x: a, y: _ } = Position{ x: 2, y: 3 };
// unused result, assignment to `_` used to declare intent and
remove a warning
_ = 2 + 2;
// triggers unused_must_use warning
// 2 + 2;
// equivalent technique using a wildcard pattern in a let-
binding
// 2 + 2;
```

let \_ = 2 + 2;

r[patterns]

# **Patterns**

r[patterns.syntax]

Pattern -> `|`? PatternNoTopAlt ( `|` PatternNoTopAlt )\*

PatternNoTopAlt ->

PatternWithoutRange

| RangePattern

PatternWithoutRange ->

- LiteralPattern
- | IdentifierPattern
- | WildcardPattern
- | RestPattern
- | ReferencePattern
- | StructPattern
- | TupleStructPattern
- | TuplePattern
- | GroupedPattern
- | SlicePattern
- | PathPattern
- | MacroInvocation

r[patterns.intro] Patterns are used to match values against structures and to, optionally, bind variables to values inside these structures. They are also used in variable declarations and parameters for functions and closures.

The pattern in the following example does four things:

- Tests if person has the car field filled with something.
- Tests if the person's age field is between 13 and 19, and binds its value to the person\_age variable.
- Binds a reference to the name field to the variable person\_name.
- Ignores the rest of the fields of person. The remaining fields can have any value and are not bound to any variables.

```
# struct Car;
# struct Computer;
# struct Person {
# name: String,
# car: Option<Car>,
     computer: Option<Computer>,
#
#
      age: u8,
# }
# let person = Person {
      name: String::from("John"),
#
     car: Some(Car),
#
      computer: None,
#
#
      age: 15,
# };
if let
    Person {
        car: Some(_),
        age: person_age @ 13..=19,
        name: ref person_name,
        . .
    } = person
{
    println!("{} has a car and is {} years old.", person_name,
person_age);
}
```

```
r[patterns.usage] Patterns are used in:
r[patterns.let]
```

```
• <u>let declarations</u>
```

```
r[patterns.param]
```

- <u>Function</u> and <u>closure</u> parameters r[patterns.match]
- match expressions

r[patterns.if-let]

- <u>if let expressions</u> r[patterns.while-let]
- while let expressions
  r[patterns.for]
- <u>for expressions</u> r[patterns.destructure]

### Destructuring

r[patterns.destructure.intro] Patterns can be used to *destructure* <u>structs</u>, <u>enums</u>, and <u>tuples</u>. Destructuring breaks up a value into its component pieces. The syntax used is almost the same as when creating such values.

r[patterns.destructure.wildcard] In a pattern whose <u>scrutinee</u> expression has a <u>struct</u>, <u>enum</u> or <u>tuple</u> type, a <u>wildcard pattern</u> (\_) stands in for a *single* data field, whereas an <u>et cetera</u> or <u>rest pattern</u> (..) stands in for *all* the remaining fields of a particular variant.

r[patterns.destructure.named-field-shorthand] When destructuring a data structure with named (but not numbered) fields, it is allowed to write fieldname as a shorthand for fieldname: fieldname.

```
# enum Message {
#
      Quit,
#
      WriteString(String),
      Move { x: i32, y: i32 },
#
      ChangeColor(u8, u8, u8),
#
# }
# let message = Message::Quit;
match message {
    Message::Quit => println!("Quit"),
    Message::WriteString(write) => println!("{}", &write),
        Message::Move{ x, y: 0 } => println!("move
                                                             {}
horizontally", x),
    Message::Move{ .. } => println!("other move"),
    Message::ChangeColor { 0: red, 1: green, 2: _ } => {
           println!("color change, red: {}, green: {}", red,
green);
    }
};
```

r[patterns.refutable]

#### Refutability

A pattern is said to be *refutable* when it has the possibility of not being matched by the value it is being matched against. *Irrefutable* patterns, on the other hand, always match the value they are being matched against. Examples:

r[patterns.literal]

## **Literal patterns**

```
r[patterns.literal.syntax]
LiteralPattern -> `-`? LiteralExpression
```

r[patterns.literal.intro] *Literal patterns* match exactly the same value as what is created by the literal. Since negative numbers are not <u>literals</u>, literals in patterns may be prefixed by an optional minus sign, which acts like the negation operator.

[!WARNING] C string and raw C string literals are accepted in literal patterns, but &CStr doesn't implement structural equality (# [derive(Eq, PartialEq)]) and therefore any such match on a &CStr will be rejected with a type error.

r[patterns.literal.refutable] Literal patterns are always refutable.

Examples:

```
for i in -2..5 {
    match i {
        -1 => println!("It's minus one"),
        1 => println!("It's a one"),
        2|4 => println!("It's either a two or a four"),
        _ => println!("Matched none of the arms"),
    }
}
```

r[patterns.ident]

## **Identifier patterns**

r[patterns.ident.syntax]

IdentifierPattern -> `ref`? `mut`? IDENTIFIER ( `@` PatternNoTopAlt )?

r[patterns.ident.intro] Identifier patterns bind the value they match to a variable in the <u>value namespace</u>.

r[patterns.ident.unique] The identifier must be unique within the pattern.

r[patterns.ident.scope] The variable will shadow any variables of the same name in scope. The <u>scope</u> of the new binding depends on the context of where the pattern is used (such as a let binding or a match arm).

r[patterns.ident.bare] Patterns that consist of only an identifier, possibly with a mut, match any value and bind it to that identifier. This is the most commonly used pattern in variable declarations and parameters for functions and closures.

r[patterns.ident.scrutinized] To bind the matched value of a pattern to a variable, use the syntax variable @ subpattern. For example, the following binds the value 2 to e (not the entire range: the range here is a range subpattern).

```
let x = 2;
match x {
    e @ 1 ..= 5 => println!("got a range element {}", e),
    _ => println!("anything"),
}
```

r[patterns.ident.move] By default, identifier patterns bind a variable to a copy of or move from the matched value depending on whether the matched value implements <u>Copy</u>.

r[patterns.ident.ref] This can be changed to bind to a reference by using the ref keyword, or to a mutable reference using ref mut. For example:

```
# let a = Some(10);
match a {
    None => (),
    Some(value) => (),
}
match a {
    None => (),
    Some(ref value) => (),
}
```

In the first match expression, the value is copied (or moved). In the second match, a reference to the same memory location is bound to the variable value. This syntax is needed because in destructuring subpatterns the & operator can't be applied to the value's fields. For example, the following is not valid:

```
# struct Person {
# name: String,
# age: u8,
# }
# let value = Person { name: String::from("John"), age: 23 };
if let Person { name: &person_name, age: 18..=150 } = value { }
```

To make it valid, write the following:

```
# struct Person {
# name: String,
# age: u8,
# }
# let value = Person { name: String::from("John"), age: 23 };
if let Person { name: ref person_name, age: 18..=150 } = value
{ }
```

r[patterns.ident.ref-ignored] Thus, ref is not something that is being matched against. Its objective is exclusively to make the matched binding a reference, instead of potentially copying or moving what was matched.

r[patterns.ident.precedent] <u>Path patterns</u> take precedence over identifier patterns.

r[patterns.ident.constraint] It is an error if ref or ref mut is specified and the identifier shadows a constant.

r[patterns.ident.refutable] Identifier patterns are irrefutable if the @ subpattern is irrefutable or the subpattern is not specified.

r[patterns.ident.binding]

## **Binding modes**

r[patterns.ident.binding.intro] To service better ergonomics, patterns operate in different *binding modes* in order to make it easier to bind references to values. When a reference value is matched by a non-reference pattern, it will be automatically treated as a ref or ref mut binding. Example:

```
let x: &Option<i32> = &Some(3);
if let Some(y) = x {
    // y was converted to `ref y` and its type is &i32
}
```

r[patterns.ident.binding.non-reference] *Non-reference patterns* include all patterns except bindings, <u>wildcard patterns</u> (\_\_), <u>const patterns</u> of reference types, and <u>reference patterns</u>.

r[patterns.ident.binding.default-mode] If a binding pattern does not explicitly have ref, ref mut, or mut, then it uses the *default binding mode* to determine how the variable is bound.

r[patterns.ident.binding.move] The default binding mode starts in "move" mode which uses move semantics.

r[patterns.ident.binding.top-down] When matching a pattern, the compiler starts from the outside of the pattern and works inwards.

r[patterns.ident.binding.auto-deref] Each time a reference is matched using a non-reference pattern, it will automatically dereference the value and update the default binding mode.

r[patterns.ident.binding.ref] References will set the default binding mode to ref.

r[patterns.ident.binding.ref-mut] Mutable references will set the mode to ref mut unless the mode is already ref in which case it remains ref.

r[patterns.ident.binding.nested-references] If the automatically dereferenced value is still a reference, it is dereferenced and this process repeats.

r[patterns.ident.binding.mode-limitations-binding] The binding pattern may only explicitly specify a ref or ref mut binding mode, or specify mutability with mut, when the default binding mode is "move". For example, these are not accepted:

let [mut x] = &[()]; //~ ERROR
let [ref x] = &[()]; //~ ERROR
let [ref mut x] = &mut [()]; //~ ERROR

r[patterns.ident.binding.mode-limitations.edition2024]

[!EDITION-2024] Before the 2024 edition, bindings could explicitly specify a ref or ref mut binding mode even when the default binding mode was not "move", and they could specify mutability on such bindings with mut. In these editions, specifying mut on a binding set the binding mode to "move" regardless of the current default binding mode.

r[patterns.ident.binding.mode-limitations-reference] Similarly, a reference pattern may only appear when the default binding mode is "move". For example, this is not accepted:

let [&x] = &[&()]; //~ ERROR

r[patterns.ident.binding.mode-limitations-reference.edition2024]

[!EDITION-2024] Before the 2024 edition, reference patterns could appear even when the default binding mode was not "move", and had both the effect of matching against the scrutinee and of causing the default binding mode to be reset to "move".

r[patterns.ident.binding.mixed] Move bindings and reference bindings can be mixed together in the same pattern. Doing so will result in partial move of the object bound to and the object cannot be used afterwards. This applies only if the type cannot be copied.

In the example below, name is moved out of person. Trying to use person as a whole or person.name would result in an error because of *partial move*.

Example:

```
# struct Person {
# name: String,
# age: u8,
# }
# let person = Person{ name: String::from("John"), age: 23 };
// `name` is moved from person and `age` referenced
let Person { name, ref age } = person;
```

r[patterns.wildcard]

## Wildcard pattern

```
r[patterns.wildcard.syntax]
WildcardPattern -> `_`
```

r[patterns.wildcard.intro] The *wildcard pattern* (an underscore symbol) matches any value. It is used to ignore values when they don't matter.

r[patterns.wildcard.struct-matcher] Inside other patterns it matches a single data field (as opposed to the ... which matches the remaining fields).

r[patterns.wildcard.no-binding] Unlike identifier patterns, it does not copy, move or borrow the value it matches.

Examples:

```
# let x = 20;
let (a, \_) = (10, x); // the x is always matched by _
# assert_eq!(a, 10);
// ignore a function/closure param
let real_part = |a: f64, _: f64| { a };
// ignore a field from a struct
# struct RGBA {
#
     r: f32,
#
     g: f32,
     b: f32,
#
#
     a: f32,
# }
# let color = RGBA{r: 0.4, g: 0.1, b: 0.9, a: 0.5};
let RGBA{r: red, g: green, b: blue, a: _} = color;
# assert_eq!(color.r, red);
# assert_eq!(color.g, green);
# assert_eq!(color.b, blue);
// accept any Some, with any value
# let x = Some(10);
if let Some(_) = x {}
```

r[patterns.wildcard.refutable] The wildcard pattern is always irrefutable. r[patterns.rest]

#### **Rest patterns**

```
r[patterns.rest.syntax]
RestPattern -> `..`
```

r[patterns.rest.intro] The *rest pattern* (the . . token) acts as a variablelength pattern which matches zero or more elements that haven't been matched already before and after.

r[patterns.rest.allowed-patterns] It may only be used in <u>tuple</u>, <u>tuple</u> <u>struct</u>, and <u>slice</u> patterns, and may only appear once as one of the elements in those patterns. It is also allowed in an <u>identifier pattern</u> for <u>slice patterns</u> only.

r[patterns.rest.refutable] The rest pattern is always irrefutable.

Examples:

```
# let words = vec!["a", "b", "c"];
# let slice = &words[..];
match slice {
    [] => println!("slice is empty"),
    [one] => println!("single element {}", one),
    [head, tail @ ..] => println!("head={} tail={:?}", head,
tail),
}
match slice {
    // Ignore everything but the last element, which must be
"!".
    [.., "!"] => println!("!!!"),
      // `start` is a slice of everything except the last
element, which must be "z".
    [start @ .., "z"] => println!("starts with: {:?}", start),
    // `end` is a slice of everything but the first element,
which must be "a".
    ["a", end @ ..] => println!("ends with: {:?}", end),
```

```
// 'whole' is the entire slice and `last` is the final
element
    whole @ [.., last] => println!("the last element of {:?}
is {}", whole, last),
    rest => println!("{:?}", rest),
}
if let [.., penultimate, _] = slice {
    println!("next to last is {}", penultimate);
}
\# let tuple = (1, 2, 3, 4, 5);
// Rest patterns may also be used in tuple and tuple struct
patterns.
match tuple {
    (1, .., y, z) => println!("y={} z={}", y, z),
    (.., 5) => println!("tail must be 5"),
    (..) => println!("matches everything else"),
}
```

r[patterns.range]

## **Range patterns**

```
r[patterns.range.syntax]
RangePattern ->
      RangeExclusivePattern
    | RangeInclusivePattern
    | RangeFromPattern
    | RangeToExclusivePattern
    | RangeToInclusivePattern
    ObsoleteRangePattern[^obsolete-range-edition]
RangeExclusivePattern ->
      RangePatternBound `..` RangePatternBound
RangeInclusivePattern ->
      RangePatternBound `..=` RangePatternBound
RangeFromPattern ->
      RangePatternBound `..`
RangeToExclusivePattern ->
      `..` RangePatternBound
RangeToInclusivePattern ->
      `..=` RangePatternBound
ObsoleteRangePattern ->
    RangePatternBound `...` RangePatternBound
RangePatternBound ->
      LiteralPattern
    | PathExpression
1
```

The [ObsoleteRangePattern] syntax has been removed in the 2021 edition.

r[patterns.range.intro] *Range patterns* match scalar values within the range defined by their bounds. They comprise a *sigil* ( . . or . . = ) and a bound on one or both sides.

A bound on the left of the sigil is called a *lower bound*. A bound on the right is called an *upper bound*.

r[patterns.range.exclusive] The *exclusive range pattern* matches all values from the lower bound up to, but not including the upper bound. It is written as its lower bound, followed by ..., followed by the upper bound.

For example, a pattern 'm'..'p' will match only 'm', 'n' and 'o', specifically **not** including 'p'.

r[patterns.range.inclusive] The *inclusive range pattern* matches all values from the lower bound up to and including the upper bound. It is written as its lower bound, followed by ...=, followed by the upper bound.

For example, a pattern 'm'..='p' will match only the values 'm', 'n', 'o', and 'p'.

r[patterns.range.from] The *from range pattern* matches all values greater than or equal to the lower bound. It is written as its lower bound followed by . . .

For example, **1**.. will match any integer greater than or equal to 1, such as 1, 9, or 9001, or 9007199254740991 (if it is of an appropriate size), but not 0, and not negative numbers for signed integers.

r[patterns.range.to-exclusive] The *to exclusive range pattern* matches all values less than the upper bound. It is written as . followed by the upper bound.

For example, ...10 will match any integer less than 10, such as 9, 1, 0, and for signed integer types, all negative values.

r[patterns.range.to-inclusive] The *to inclusive range pattern* matches all values less than or equal to the upper bound. It is written as ..= followed by the upper bound.

For example, ..=10 will match any integer less than or equal to 10, such as 10, 1, 0, and for signed integer types, all negative values.

r[patterns.range.constraint-less-than] The lower bound cannot be greater than the upper bound. That is, in a..=b,  $a \le b$  must be the case. For example, it is an error to have a range pattern 10..=0.

r[patterns.range.bound] A bound is written as one of:

- A character, byte, integer, or float literal.
- A followed by an integer or float literal.
- A <u>path</u>.

[!NOTE]

We syntactically accept more than this for a *[RangePatternBound]*. We later reject the other things semantically.

r[patterns.range.constraint-bound-path] If a bound is written as a path, after macro resolution, the path must resolve to a constant item of the type char, an integer type, or a float type.

r[patterns.range.type] The range pattern matches the type of its upper and lower bounds, which must be the same type.

r[patterns.range.path-value] If a bound is a <u>path</u>, the bound matches the type and has the value of the <u>constant</u> the path resolves to.

r[patterns.range.literal-value] If a bound is a literal, the bound matches the type and has the value of the corresponding <u>literal expression</u>.

r[patterns.range.negation] If a bound is a literal preceded by a -, the bound matches the same type as the corresponding <u>literal expression</u> and has the value of <u>negating</u> the value of the corresponding literal expression.

r[patterns.range.float-restriction] For float range patterns, the constant may not be a NaN.

**Examples:** 

```
# let c = 'f';
let valid_variable = match c {
    'a'..='z' => true,
    'A'..='Z' => true,
    'a'..='\outletu' => true,
    _ => false,
```

};

```
# let ph = 10;
println!("{}", match ph {
    0..7 => "acid",
    7 \Rightarrow "neutral",
    8..=14 => "base",
    _ => unreachable!(),
});
# let uint: u32 = 5;
match uint {
    0 => "zero!",
    1.. => "positive number!",
};
// using paths to constants:
# const TROPOSPHERE MIN : u8 = 6;
# const TROPOSPHERE MAX : u8 = 20;
#
# const STRATOSPHERE_MIN : u8 = TROPOSPHERE_MAX + 1;
# const STRATOSPHERE MAX : u8 = 50;
#
# const MESOSPHERE_MIN : u8 = STRATOSPHERE_MAX + 1;
# const MESOSPHERE_MAX : u8 = 85;
#
# let altitude = 70;
#
println!("{}", match altitude {
    TROPOSPHERE_MIN..=TROPOSPHERE_MAX => "troposphere",
    STRATOSPHERE_MIN..=STRATOSPHERE_MAX => "stratosphere",
    MESOSPHERE_MIN..=MESOSPHERE_MAX => "mesosphere",
    _ => "outer space, maybe",
});
# pub mod binary {
```

```
#
      pub const MEGA : u64 = 1024*1024;
#
      pub const GIGA : u64 = 1024*1024*1024;
# }
# let n_items = 20_832_425;
# let bytes_per_item = 12;
if let size @
                   binary::MEGA..=binary::GIGA = n_items
bytes_per_item {
    println!("It fits and occupies {} bytes", size);
}
# trait MaxValue {
      const MAX: u64;
#
# }
# impl MaxValue for u8 {
      const MAX: u64 = (1 << 8) - 1;
#
# }
# impl MaxValue for u16 {
      const MAX: u64 = (1 << 16) - 1;
#
# }
# impl MaxValue for u32 {
      const MAX: u64 = (1 << 32) - 1;
#
# }
// using qualified paths:
println!("{}", match 0xfacade {
    0 ..= <u8 as MaxValue>::MAX => "fits in a u8",
    0 ..= <u16 as MaxValue>::MAX => "fits in a u16",
    0 ..= <u32 as MaxValue>::MAX => "fits in a u32",
    _ => "too big",
});
```

r[patterns.range.refutable] Range patterns for fix-width integer and char types are irrefutable when they span the entire set of possible values of a type. For example, 0u8..=255u8 is irrefutable.

r[patterns.range.refutable-integer] The range of values for an integer type is the closed range from its minimum to maximum value.

r[patterns.range.refutable-char] The range of values for a char type are
precisely those ranges containing all Unicode Scalar Values:
'\u{0000}'..='\u{D7FF}' and '\u{E000}'..='\u{10FFFF}'.

r[patterns.range.constraint-slice] [RangeFromPattern] cannot be used as a top-level pattern for subpatterns in <u>slice patterns</u>. For example, the pattern [1..., \_] is not a valid pattern.

r[patterns.range.edition2021]

[!EDITION-2021] Before the 2021 edition, range patterns with both a lower and upper bound may also be written using ... in place of ..=, with the same meaning.

r[patterns.ref]

#### **Reference patterns**

r[patterns.ref.syntax]

```
ReferencePattern -> (`&`|`&&`) `mut`? PatternWithoutRange
```

r[patterns.ref.intro] Reference patterns dereference the pointers that are being matched and, thus, borrow them.

For example, these two matches on X: &i32 are equivalent:

```
let int_reference = &3;
let a = match *int_reference { 0 => "zero", _ => "some" };
let b = match int_reference { &0 => "zero", _ => "some" };
```

```
assert_eq!(a, b);
```

r[patterns.ref.ref-ref] The grammar production for reference patterns has to match the token && to match a reference to a reference because it is a token by itself, not two & tokens.

r[patterns.ref.mut] Adding the mut keyword dereferences a mutable reference. The mutability must match the mutability of the reference.

r[patterns.ref.refutable] Reference patterns are always irrefutable.
r[patterns.struct]

#### **Struct patterns**

```
r[patterns.struct.syntax]
StructPattern ->
    PathInExpression `{`
        StructPatternElements?
    `}`
StructPatternElements ->
      StructPatternFields (`,` | `,` StructPatternEtCetera)?
    | StructPatternEtCetera
StructPatternFields ->
    StructPatternField (`,` StructPatternField)*
StructPatternField ->
    OuterAttribute*
    (
        TUPLE_INDEX `:` Pattern
      | IDENTIFIER `:` Pattern
      | `ref`? `mut`? IDENTIFIER
    )
```

```
StructPatternEtCetera -> `..`
```

r[patterns.struct.intro] Struct patterns match struct, enum, and union values that match all criteria defined by its subpatterns. They are also used to <u>destructure</u> a struct, enum, or union value.

r[patterns.struct.ignore-rest] On a struct pattern, the fields are referenced by name, index (in the case of tuple structs) or ignored by use of ...:

```
#
match s {
    Point {x: 10, y: 20} => (),
    Point {y: 10, x: 20} => (), // order doesn't matter
    Point {x: 10, ..} => (),
    Point {..} => (),
}
# struct PointTuple (
#
      u32,
#
      u32,
# );
# let t = PointTuple(1, 2);
#
match t {
    PointTuple {0: 10, 1: 20} => (),
    PointTuple {1: 10, 0: 20} => (), // order doesn't matter
    PointTuple {0: 10, ..} => (),
    PointTuple \{..\} => (),
}
# enum Message {
#
      Quit,
#
      Move { x: i32, y: i32 },
# }
# let m = Message::Quit;
#
match m {
    Message::Quit => (),
    Message::Move {x: 10, y: 20} => (),
    Message::Move \{..\} => (),
}
```

r[patterns.struct.constraint-struct] If . . is not used, a struct pattern used to match a struct is required to specify all fields:

```
# struct Struct {
#
     a: i32,
     b: char,
#
#
     c: bool,
# }
# let mut struct_value = Struct{a: 10, b: 'X', c: false};
#
match struct_value {
    Struct{a: 10, b: 'X', c: false} => (),
    Struct{a: 10, b: 'X', ref c} => (),
    Struct{a: 10, b: 'X', ref mut c} => (),
    Struct{a: 10, b: 'X', c: _} => (),
    Struct{a: _, b: _, c: _} => (),
}
```

r[patterns.struct.constraint-union] A struct pattern used to match a union must specify exactly one field (see <u>Pattern matching on unions</u>).

r[patterns.struct.binding-shorthand] The ref and/or mut [IDENTIFIER] syntax matches any value and binds it to a variable with the same name as the given field.

r[patterns.struct.refutable] A struct pattern is refutable if the [PathInExpression] resolves to a constructor of an enum with more than one variant, or one of its subpatterns is refutable.

r[patterns.struct.namespace] A struct pattern matches against the struct, union, or enum variant whose constructor is resolved from

[PathInExpression] in the <u>type namespace</u>. See [patterns.tuple-struct.namespace] for more details.

r[patterns.tuple-struct]

## **Tuple struct patterns**

r[patterns.tuple-struct.syntax]

TupleStructPattern -> PathInExpression `(` TupleStructItems?
`)`

```
TupleStructItems -> Pattern ( `,` Pattern )* `,`?
```

r[patterns.tuple-struct.intro] Tuple struct patterns match tuple struct and enum values that match all criteria defined by its subpatterns. They are also used to <u>destructure</u> a tuple struct or enum value.

r[patterns.tuple-struct.refutable] A tuple struct pattern is refutable if the [PathInExpression] resolves to a constructor of an enum with more than one variant, or one of its subpatterns is refutable.

r[patterns.tuple-struct.namespace] A tuple struct pattern matches against the tuple struct or <u>tuple-like enum variant</u> whose constructor is resolved from [PathInExpression] in the <u>value namespace</u>.

```
[!NOTE] Conversely, a struct pattern for a tuple struct or <u>tuple-like</u>
<u>enum variant</u>, e.g. S { 0: _ }, matches against the tuple struct or
variant whose constructor is resolved in the <u>type namespace</u>.
enum E1 { V(u16) }
enum E2 { V(u32) }
// Import `E1::V` from the type namespace only.
mod _0 {
    const V: () = (); // For namespace masking.
    pub(super) use super::E1::*;
}
use _0::*;
// Import `E2::V` from the value namespace only.
mod 1 \{
    struct V {} // For namespace masking.
    pub(super) use super::E2::*;
}
```

```
use _1::*;
fn f() {
    // This struct pattern matches against the tuple-like
    // enum variant whose constructor was found in the type
    // namespace.
    let V { 0: ..=u16::MAX } = (loop {}) else { loop {} };
    // This tuple struct pattern matches against the tuple-
like
     // enum variant whose constructor was found in the
value
    // namespace.
    let V(..=u32::MAX) = (loop {}) else { loop {} };
}
# // Required due to the odd behavior of `super` within
functions.
# fn main() {}
```

The Lang team has made certain decisions, such as in <u>PR #138458</u>, that raise questions about the desirability of using the value namespace in this way for patterns, as described in <u>PR #140593</u>. It might be prudent to not intentionally rely on this nuance in your code.

```
r[patterns.tuple]
```

#### **Tuple patterns**

```
r[patterns.tuple.syntax]
TuplePattern -> `(` TuplePatternItems? `)`
```

```
TuplePatternItems ->
Pattern `,`
| RestPattern
| Pattern (`,` Pattern)+ `,`?
```

r[patterns.tuple.intro] Tuple patterns match tuple values that match all criteria defined by its subpatterns. They are also used to <u>destructure</u> a tuple.

r[patterns.tuple.rest-syntax] The form (...) with a single [RestPattern] is a special form that does not require a comma, and matches a tuple of any size.

r[patterns.tuple.refutable] The tuple pattern is refutable when one of its subpatterns is refutable.

An example of using tuple patterns:

```
let pair = (10, "ten");
let (a, b) = pair;
assert_eq!(a, 10);
assert_eq!(b, "ten");
```

r[patterns.paren]

## **Grouped patterns**

```
r[patterns.paren.syntax]
GroupedPattern -> `(` Pattern `)`
```

r[patterns.paren.intro] Enclosing a pattern in parentheses can be used to explicitly control the precedence of compound patterns. For example, a reference pattern next to a range pattern such as &0..=5 is ambiguous and is not allowed, but can be expressed with parentheses.

```
let int_reference = &3;
match int_reference {
    &(0..=5) => (),
    _ => (),
}
```

r[patterns.slice]

#### **Slice patterns**

```
r[patterns.slice.syntax]
SlicePattern -> `[` SlicePatternItems? `]`
```

```
SlicePatternItems -> Pattern (`,` Pattern)* `,`?
```

r[patterns.slice.intro] Slice patterns can match both arrays of fixed size and slices of dynamic size.

```
// Fixed size
let arr = [1, 2, 3];
match arr {
    [1, _, _] => "starts with one",
    [a, b, c] => "starts with something else",
};
// Dynamic size
let v = vec![1, 2, 3];
match v[..] {
    [a, b] => { /* this arm will not apply because the length
doesn't match */ }
    [a, b, c] => { /* this arm will apply */ }
    _ => { /* this wildcard is required, since the length is
not known statically */ }
};
```

r[patterns.slice.refutable-array] Slice patterns are irrefutable when matching an array as long as each element is irrefutable.

r[patterns.slice.refutable-slice] When matching a slice, it is irrefutable only in the form with a single . . <u>rest pattern</u> or <u>identifier pattern</u> with the . . rest pattern as a subpattern.

r[patterns.slice.restriction] Within a slice, a range pattern without both lower and upper bound must be enclosed in parentheses, as in (a..), to clarify it is intended to match against a single slice element. A range pattern with both lower and upper bound, like a..=b, is not required to be enclosed in parentheses.

r[patterns.path]

## **Path patterns**

r[patterns.path.syntax]
PathPattern -> PathExpression

r[patterns.path.intro] *Path patterns* are patterns that refer either to constant values or to structs or enum variants that have no fields.

r[patterns.path.unqualified] Unqualified path patterns can refer to:

- enum variants
- structs
- constants
- associated constants

r[patterns.path.qualified] Qualified path patterns can only refer to associated constants.

r[patterns.path.refutable] Path patterns are irrefutable when they refer to structs or an enum variant when the enum has only one variant or a constant whose type is irrefutable. They are refutable when they refer to refutable constants or enum variants for enums with multiple variants.

r[patterns.const]

#### **Constant patterns**

r[patterns.const.partial-eq] When a constant C of type T is used as a pattern, we first check that T: PartialEq.

r[patterns.const.structural-equality] Furthermore we require that the value of **C** *has (recursive) structural equality*, which is defined recursively as follows:

r[patterns.const.primitive]

• Integers as well as str, bool and char values always have structural equality.

r[patterns.const.builtin-aggregate]

• Tuples, arrays, and slices have structural equality if all their fields/elements have structural equality. (In particular, () and []

always have structural equality.) r[patterns.const.ref]

• References have structural equality if the value they point to has structural equality.

r[patterns.const.aggregate]

- A value of struct or enum type has structural equality if its
   PartialEq instance is derived via #[derive(PartialEq)], and all
   fields (for enums: of the active variant) have structural equality.

   r[patterns.const.pointer]
- A raw pointer has structural equality if it was defined as a constant integer (and then cast/transmuted).

r[patterns.const.float]

• A float value has structural equality if it is not a NaN.

r[patterns.const.exhaustive]

• Nothing else has structural equality.

r[patterns.const.generic] In particular, the value of **c** must be known at pattern-building time (which is pre-monomorphization). This means that associated consts that involve generic parameters cannot be used as patterns.

r[patterns.const.immutable] The value of C must not contain any references to mutable statics (static mut items or interior mutable static items) or extern statics.

r[patterns.const.translation] After ensuring all conditions are met, the constant value is translated into a pattern, and now behaves exactly as-if that pattern had been written directly. In particular, it fully participates in exhaustiveness checking. (For raw pointers, constants are the only way to write such patterns. Only \_\_\_\_\_ is ever considered exhaustive for these types.)

r[patterns.or]
#### **Or-patterns**

*Or-patterns* are patterns that match on one of two or more sub-patterns (for example  $A \mid B \mid C$ ). They can nest arbitrarily. Syntactically, or-patterns are allowed in any of the places where other patterns are allowed (represented by the [Pattern] production), with the exceptions of let-bindings and function and closure arguments (represented by the [PatternNoTopAlt] production).

r[patterns.constraints]

#### **Static semantics**

r[patterns.constraints.pattern]

- 1. Given a pattern **p** | **q** at some depth for some arbitrary patterns **p** and **q**, the pattern is considered ill-formed if:
  - the type inferred for p does not unify with the type inferred for q, or
  - the same set of bindings are not introduced in p and q, or
  - the type of any two bindings with the same name in p and q do not unify with respect to types or binding modes.

Unification of types is in all instances aforementioned exact and implicit <u>type coercions</u> do not apply.

r[patterns.constraints.match-type-check] 2. When type checking an expression match e\_s { a\_1 => e\_1, ... a\_n => e\_n }, for each match arm a\_i which contains a pattern of form p\_i | q\_i, the pattern p\_i | q\_i is considered ill formed if, at the depth d where it exists the fragment of e\_s at depth d, the type of the expression fragment does not unify with p\_i | q\_i.

r[patterns.constraints.exhaustiveness-or-pattern] 3. With respect to exhaustiveness checking, a pattern  $p \mid q$  is considered to cover p as well as q. For some constructor c(x, ...) the distributive law applies such that  $c(p \mid q, ...rest)$  covers the same set of value as  $c(p, ...rest) \mid c(q, ...rest)$ 

..rest) does. This can be applied recursively until there are no more nested patterns of form  $p \mid q$  other than those that exist at the top level.

Note that by *"constructor"* we do not refer to tuple struct patterns, but rather we refer to a pattern for any product type. This includes enum variants, tuple structs, structs with named fields, arrays, tuples, and slices.

r[patterns.behavior]

#### **Dynamic semantics**

r[patterns.behavior.nested-or-patterns]

1. The dynamic semantics of pattern matching a scrutinee expression e\_s against a pattern c(p | q, ...rest) at depth d where c is some constructor, p and q are arbitrary patterns, and rest is optionally any remaining potential factors in c, is defined as being the same as that of c(p, ...rest) | c(q, ...rest).

r[patterns.precedence]

#### **Precedence** with other undelimited patterns

As shown elsewhere in this chapter, there are several types of patterns that are syntactically undelimited, including identifier patterns, reference patterns, and or-patterns. Or-patterns always have the lowest-precedence. This allows us to reserve syntactic space for a possible future type ascription feature and also to reduce ambiguity. For example,  $\times @ A(..) | B(..)$  will result in an error that  $\times$  is not bound in all patterns.  $\&A(\times) | B(\times)$  will result in a type mismatch between  $\times$  in the different subpatterns.

# Type system

r[type]

# **Types**

r[type.intro] Every variable, item, and value in a Rust program has a type. The *type* of a *value* defines the interpretation of the memory holding it and the operations that may be performed on the value.

r[type.builtin] Built-in types are tightly integrated into the language, in nontrivial ways that are not possible to emulate in user-defined types.

r[type.user-defined] User-defined types have limited capabilities.

r[type.kinds] The list of types is:

- Primitive types:
  - <u>Boolean</u> --- bool
  - <u>Numeric</u> --- integer and float
  - Textual --- char and str
  - <u>Never</u> --- ! --- a type with no values
- Sequence types:
  - <u>Tuple</u>
  - <u>Array</u>
  - <u>Slice</u>
- User-defined types:
  - <u>Struct</u>
  - <u>Enum</u>
  - <u>Union</u>
- Function types:
  - <u>Functions</u>
  - <u>Closures</u>
- Pointer types:
  - <u>References</u>
  - <u>Raw pointers</u>
  - Function pointers

- Trait types:
  - <u>Trait objects</u>
  - <u>Impl trait</u>

r[type.name]

### **Type expressions**

r[type.name.syntax]

Туре ->

TypeNoBounds

- | ImplTraitType
- | TraitObjectType

TypeNoBounds ->

ParenthesizedType

| ImplTraitTypeOneBound

| TraitObjectTypeOneBound

| TypePath

- | ТирlеТуре
- | NeverType
- | RawPointerType
- | ReferenceType
- | ArrayType
- | SliceType
- | InferredType
- | QualifiedPathInType
- | BareFunctionType
- | MacroInvocation

r[type.name.intro] A *type expression* as defined in the [Type] grammar rule above is the syntax for referring to a type. It may refer to:

r[type.name.sequence]

- Sequence types (<u>tuple</u>, <u>array</u>, <u>slice</u>). r[type.name.path]
- <u>Type paths</u> which can reference:
  - Primitive types (<u>boolean</u>, <u>numeric</u>, <u>textual</u>).
  - Paths to an <u>item (struct, enum, union, type alias</u>, <u>trait</u>).
  - <u>Self path</u> where Self is the implementing type.

• Generic <u>type parameters</u>.

r[type.name.pointer]

• Pointer types (<u>reference</u>, <u>raw pointer</u>, <u>function pointer</u>). r[type.name.inference]

• The <u>inferred type</u> which asks the compiler to determine the type. r[type.name.grouped]

• <u>Parentheses</u> which are used for disambiguation. r[type.name.trait]

• Trait types: <u>Trait objects</u> and <u>impl trait</u>. r[type.name.never]

• The <u>never</u> type. r[type.name.macro-expansion]

• <u>Macros</u> which expand to a type expression. r[type.name.parenthesized]

### **Parenthesized types**

```
r[type.name.parenthesized.syntax]
ParenthesizedType -> `(` Type `)`
```

r[type.name.parenthesized.intro] In some situations the combination of types may be ambiguous. Use parentheses around a type to avoid ambiguity. For example, the + operator for <u>type boundaries</u> within a <u>reference type</u> is unclear where the boundary applies, so the use of parentheses is required. Grammar rules that require this disambiguation use the [TypeNoBounds] rule instead of [Type][grammar-Type].

```
# use std::any::Any;
type T<'a> = &'a (dyn Any + Send);
```

r[type.recursive]

#### **Recursive types**

r[type.recursive.intro] Nominal types — <u>structs</u>, <u>enumerations</u>, and <u>unions</u> — may be recursive. That is, each <u>enum</u> variant or <u>struct</u> or <u>union</u> field may refer, directly or indirectly, to the enclosing <u>enum</u> or <u>struct</u> type itself.

r[type.recursive.constraint] Such recursion has restrictions:

- Recursive types must include a nominal type in the recursion (not mere type aliases, or other structural types such as arrays or tuples). So type Rec = &'static [Rec] is not allowed.
- The size of a recursive type must be finite; in other words the recursive fields of the type must be <u>pointer types</u>.

An example of a *recursive* type and its use:

```
enum List<T> {
    Nil,
    Cons(T, Box<List<T>>)
}
let a: List<i32> = List::Cons(7, Box::new(List::Cons(13,
Box::new(List::Nil))));
```

r[type.bool]

## **Boolean type**

```
let b: bool = true;
```

r[type.bool.intro] The *boolean type* or *bool* is a primitive data type that can take on one of two values, called *true* and *false*.

r[type.bool.literal] Values of this type may be created using a <u>literal</u> <u>expression</u> using the keywords true and false corresponding to the value of the same name.

r[type.bool.namespace] This type is a part of the <u>language prelude</u> with the <u>name</u> bool.

r[type.bool.layout] An object with the boolean type has a <u>size and</u> <u>alignment</u> of 1 each.

r[type.bool.repr] The value false has the bit pattern  $0 \times 00$  and the value true has the bit pattern  $0 \times 01$ . It is <u>undefined behavior</u> for an object with the boolean type to have any other bit pattern.

r[type.bool.usage] The boolean type is the type of many operands in various <u>expressions</u>:

r[type.bool.usage-condition]

• The condition operand in <u>if expressions</u> and <u>while expressions</u> r[type.bool.usage-lazy-operator]

• The operands in <u>lazy boolean operator expressions</u>

[!NOTE] The boolean type acts similarly to but is not an <u>enumerated type</u>. In practice, this mostly means that constructors are not associated to the type (e.g. bool::true).

r[type.bool.traits] Like all primitives, the boolean type <u>implements</u> the <u>traits</u> <u>Clone</u>, <u>Copy</u>, <u>Sized</u>, <u>Send</u>, and <u>Sync</u>.

[!NOTE] See the <u>standard library docs</u> for library operations.

r[type.bool.expr]

#### **Operations on boolean values**

When using certain operator expressions with a boolean type for its operands, they evaluate using the rules of <u>boolean logic</u>.

r[type.bool.expr.not]

### Logical not

b	<u>!b</u>
true	false
false	true

r[type.bool.expr.or]

## Logical or

a	b	<u>a   b</u>
true	true	true
true	false	true
false	true	true
false	false	false

r[type.bool.expr.and]

## Logical and

a	b	<u>a &amp; b</u>
true	true	true
true	false	false
false	true	false
false	false	false

r[type.bool.expr.xor]

## Logical xor

a	b	<u>a ^ b</u>
true	true	false
true	false	true
false	true	true
false	false	false

r[type.bool.expr.cmp]

## Comparisons

r[type.bool.expr.cmp.eq]

a	b	<u>a == b</u>
true	true	true
true	false	false
false	true	false
false	false	true

r[type.bool.expr.cmp.greater]

a	b	<u>a &gt; b</u>
true	true	false
true	false	true
false	true	false
false	false	false

r[type.bool.expr.cmp.not-eq]

- a != b is the same as !(a == b)
   r[type.bool.expr.cmp.greater-eq]
- a >= b is the same as a == b | a > b
   r[type.bool.expr.cmp.less]
- a < b is the same as !(a >= b)
   r[type.bool.expr.cmp.less-eq]
- a <= b is the same as a == b | a < b</li>
  r[type.bool.validity]

#### **Bit validity**

The single byte of a bool is guaranteed to be initialized (in other words, transmute::<bool, u8>(...) is always sound -- but since some bit patterns are invalid bools, the inverse is not always sound).

r[type.numeric]

# Numeric types

r[type.numeric.int]

#### **Integer types**

r[type.numeric.int.unsigned] The unsigned integer types consist of:

Туре	Minimum	Maximum
u8	0	2 <sup>8</sup> -1
u16	0	2 <sup>16</sup> -1
u32	0	2 <sup>32</sup> -1
u64	0	2 <sup>64</sup> -1
u128	0	2 <sup>128</sup> -1

r[type.numeric.int.signed] The signed two's complement integer types consist of:

Туре	Minimum	Maximum
18	<b>-</b> (2 <sup>7</sup> )	2 <sup>7</sup> -1
i16	<b>-(</b> 2 <sup>15</sup> <b>)</b>	2 <sup>15</sup> -1
i32	<b>-</b> (2 <sup>31</sup> )	2 <sup>31</sup> -1
i64	-(2 <sup>63</sup> )	2 <sup>63</sup> -1
i128	<b>-</b> (2 <sup>127</sup> )	2 <sup>127</sup> -1

r[type.numeric.float]

## **Floating-point types**

The IEEE 754-2008 "binary32" and "binary64" floating-point types are f32 and f64, respectively.

r[type.numeric.int.size]

### **Machine-dependent integer types**

r[type.numeric.int.size.usize] The usize type is an unsigned integer type with the same number of bits as the platform's pointer type. It can represent every memory address in the process.

[!NOTE] While a usize can represent every *address*, converting a *pointer* to a usize is not necessarily a reversible operation. For more information, see the documentation for <u>type cast expressions</u>, [std::ptr], and [provenance][std::ptr#provenance] in particular.

r[type.numeric.int.size.isize] The isize type is a signed integer type with the same number of bits as the platform's pointer type. The theoretical upper bound on object and array size is the maximum isize value. This ensures that isize can be used to calculate differences between pointers into an object or array and can address every byte within an object along with one byte past the end.

r[type.numeric.int.size.minimum] usize and isize are at least 16-bits wide.

[!NOTE] Many pieces of Rust code may assume that pointers, usize, and isize are either 32-bit or 64-bit. As a consequence, 16bit pointer support is limited and may require explicit care and acknowledgment from a library to support.

r[type.numeric.validity]

### **Bit validity**

For every numeric type, T, the bit validity of T is equivalent to the bit validity of  $[u8; size_of::<T>()]$ . An uninitialized byte is not a valid u8.

r[type.text]

## **Textual types**

r[type.text.intro] The types char and str hold textual data.

r[type.text.char-value] A value of type char is a <u>Unicode scalar value</u> (i.e. a code point that is not a surrogate), represented as a 32-bit unsigned word in the 0x0000 to 0xD7FF or 0xE000 to 0x10FFFF range.

r[type.text.char-precondition] It is immediate <u>undefined behavior</u> to create a char that falls outside this range. A [char] is effectively a UCS-4 / UTF-32 string of length 1.

r[type.text.str-value] A value of type str is represented the same way as [u8], a slice of 8-bit unsigned bytes. However, the Rust standard library makes extra assumptions about str: methods working on str assume and ensure that the data in there is valid UTF-8. Calling a str method with a non-UTF-8 buffer can cause <u>undefined behavior</u> now or in the future.

r[type.text.str-unsized] Since str is a <u>dynamically sized type</u>, it can only be instantiated through a pointer type, such as <u>&str</u>.

r[type.text.layout]

### Layout and bit validity

r[type.layout.char-layout] char is guaranteed to have the same size and alignment as u32 on all platforms.

r[type.layout.char-validity] Every byte of a char is guaranteed to be initialized (in other words, transmute::<char, [u8; size\_of::<char> ()]>(...) is always sound -- but since some bit patterns are invalid char s, the inverse is not always sound). r[type.never]

## Never type

```
r[type.never.syntax]
NeverType -> `!`
```

r[type.never.intro] The never type ! is a type with no values, representing the result of computations that never complete.

r[type.never.coercion] Expressions of type ! can be coerced into any other type.

r[type.never.constraint] The ! type can **only** appear in function return types presently, indicating it is a diverging function that never returns.

```
fn foo() -> ! {
    panic!("This call never returns.");
}
unsafe extern "C" {
    pub safe fn no_return_extern_func() -> !;
}
```

r[type.tuple]

# **Tuple types**

r[type.tuple.syntax]

TupleType ->

```
`(``)`
|`(`(Type`,`)+ Type?`)`
```

r[type.tuple.intro] *Tuple types* are a family of structural types<sup>1</sup> for heterogeneous lists of other types.

The syntax for a tuple type is a parenthesized, comma-separated list of types.

r[type.tuple.restriction] 1-ary tuples require a comma after their element type to be disambiguated with a <u>parenthesized type</u>.

r[type.tuple.field-number] A tuple type has a number of fields equal to the length of the list of types. This number of fields determines the *arity* of the tuple. A tuple with n fields is called an *n*-*ary tuple*. For example, a tuple with 2 fields is a 2-ary tuple.

r[type.tuple.field-name] Fields of tuples are named using increasing numeric names matching their position in the list of types. The first field is
O. The second field is 1. And so on. The type of each field is the type of the same position in the tuple's list of types.

r[type.tuple.unit] For convenience and historical reasons, the tuple type with no fields (()) is often called *unit* or *the unit type*. Its one value is also called *unit* or *the unit value*.

Some examples of tuple types:

• () (unit)

- (i32,) (1-ary tuple)
- (f64, f64)
- (String, i32)
- (i32, String) (different type from the previous example)
- (i32, f64, Vec<String>, Option<bool>)

r[type.tuple.constructor] Values of this type are constructed using a <u>tuple</u> <u>expression</u>. Furthermore, various expressions will produce the unit value if there is no other meaningful value for it to evaluate to.

r[type.tuple.access] Tuple fields can be accessed by either a <u>tuple index</u> <u>expression</u> or <u>pattern matching</u>.

Structural types are always equivalent if their internal types are equivalent. For a nominal version of tuples, see <u>tuple structs</u>.

r[type.array]

## Array types

```
r[type.array.syntax]
```

```
ArrayType -> `[` Type `;` Expression `]`
```

r[type.array.intro] An array is a fixed-size sequence of N elements of type T. The array type is written as [T; N].

r[type.array.constraint] The size is a <u>constant expression</u> that evaluates to a <u>usize</u>.

Examples:

```
// A stack-allocated array
let array: [i32; 3] = [1, 2, 3];
```

```
// A heap-allocated array, coerced to a slice
let boxed_array: Box<[i32]> = Box::new([1, 2, 3]);
```

r[type.array.index] All elements of arrays are always initialized, and access to an array is always bounds-checked in safe methods and operators.

[!NOTE] The [Vec<T>] standard library type provides a heapallocated resizable array type. r[type.slice]

## **Slice types**

r[type.slice.syntax]

```
SliceType -> `[` Type `]`
```

r[type.slice.intro] A slice is a <u>dynamically sized type</u> representing a 'view' into a sequence of elements of type **T**. The slice type is written as [T].

r[type.slice.unsized] Slice types are generally used through pointer types. For example:

- &[T]: a 'shared slice', often just called a 'slice'. It doesn't own the data it points to; it borrows it.
- &mut [T]: a 'mutable slice'. It mutably borrows the data it points to.
- Box<[T]>: a 'boxed slice'

Examples:

```
// A heap-allocated array, coerced to a slice
let boxed_array: Box<[i32]> = Box::new([1, 2, 3]);
// A (shared) slice into an array
let slice: &[i32] = &boxed_array[..];
```

r[type.slice.safe] All elements of slices are always initialized, and access to a slice is always bounds-checked in safe methods and operators.

r[type.struct]

## **Struct types**

r[type.struct.intro] A struct *type* is a heterogeneous product of other types, called the *fields* of the type. $\frac{1}{2}$ 

r[type.struct.constructor] New instances of a struct can be constructed with a struct expression.

r[type.struct.layout] The memory layout of a struct is undefined by default to allow for compiler optimizations like field reordering, but it can be fixed with the <u>repr\_attribute</u>. In either case, fields may be given in any order in a corresponding struct *expression*; the resulting struct value will always have the same memory layout.

r[type.struct.field-visibility] The fields of a struct may be qualified by visibility modifiers, to allow access to data in a struct outside a module.

r[type.struct.tuple] A *tuple struct* type is just like a struct type, except that the fields are anonymous.

r[type.struct.unit] A *unit-like struct* type is like a struct type, except that it has no fields. The one value constructed by the associated <u>struct</u> <u>expression</u> is the only value that inhabits such a type.

1

struct types are analogous to struct types in C, the *record* types of the ML family, or the *struct* types of the Lisp family.

r[type.enum]
# **Enumerated types**

r[type.enum.intro] An *enumerated type* is a nominal, heterogeneous disjoint union type, denoted by the name of an <u>enum\_item</u>.  $\frac{1}{2}$ 

r[type.enum.declaration] An <u>enum\_item</u> declares both the type and a number of *variants*, each of which is independently named and has the syntax of a struct, tuple struct or unit-like struct.

r[type.enum.constructor] New instances of an enum can be constructed with a <u>struct expression</u>.

r[type.enum.value] Any enum value consumes as much memory as the largest variant for its corresponding enum type, as well as the size needed to store a discriminant.

r[type.enum.name] Enum types cannot be denoted *structurally* as types, but must be denoted by named reference to an <u>enum\_item</u>.

1

The enum type is analogous to a data constructor declaration in Haskell, or a *pick ADT* in Limbo.

r[type.union]

# **Union types**

r[type.union.intro] A *union type* is a nominal, heterogeneous C-like union, denoted by the name of a <u>union item</u>.

r[type.union.access] Unions have no notion of an "active field". Instead, every union access transmutes parts of the content of the union to the type of the accessed field.

r[type.union.safety] Since transmutes can cause unexpected or undefined behaviour, unsafe is required to read from a union field.

r[type.union.constraint] Union field types are also restricted to a subset of types which ensures that they never need dropping. See the <u>item</u> documentation for further details.

r[type.union.layout] The memory layout of a union is undefined by default (in particular, fields do *not* have to be at offset 0), but the # [repr(...)] attribute can be used to fix a layout.

r[type.fn-item]

# **Function item types**

r[type.fn-item.intro] When referred to, a function item, or the constructor of a tuple-like struct or enum variant, yields a zero-sized value of its *function item type*.

r[type.fn-item.unique] That type explicitly identifies the function - its name, its type arguments, and its early-bound lifetime arguments (but not its late-bound lifetime arguments, which are only assigned when the function is called) - so the value does not need to contain an actual function pointer, and no indirection is needed when the function is called.

r[type.fn-item.name] There is no syntax that directly refers to a function item type, but the compiler will display the type as something like fn(u32) -> i32 {fn\_name} in error messages.

Because the function item type explicitly identifies the function, the item types of different functions - different items, or the same item with different generics - are distinct, and mixing them will create a type error:

```
fn foo<T>() { }
let x = &mut foo::<i32>;
*x = foo::<u32>; //~ ERROR mismatched types
```

r[type.fn-item.coercion] However, there is a <u>coercion</u> from function items to <u>function pointers</u> with the same signature, which is triggered not only when a function item is used when a function pointer is directly expected, but also when different function item types with the same signature meet in different arms of the same if or match:

```
# let want_i32 = false;
# fn foo<T>() { }
// `foo_ptr_1` has function pointer type `fn()` here
let foo_ptr_1: fn() = foo::<i32>;
// ... and so does `foo_ptr_2` - this type-checks.
let foo_ptr_2 = if want_i32 {
  foo::<i32>
```

} else {
 foo::<u32>
};

r[type.fn-item.traits] All function items implement [Fn], [FnMut], [FnOnce], <u>Copy</u>, <u>Clone</u>, <u>Send</u>, and <u>Sync</u>. r[type.closure]

# **Closure types**

r[type.closure.intro] A <u>closure expression</u> produces a closure value with a unique, anonymous type that cannot be written out. A closure type is approximately equivalent to a struct which contains the captured values. For instance, the following closure:

```
#[derive(Debug)]
struct Point { x: i32, y: i32 }
struct Rectangle { left_top: Point, right_bottom: Point }
fn f<F : FnOnce() -> String> (g: F) {
    println!("{}", g());
}
let mut rect = Rectangle {
    left_top: Point { x: 1, y: 1 },
    right_bottom: Point { x: 0, y: 0 }
};
let c = || {
    rect.left_top.x += 1;
    rect.right_bottom.x += 1;
    format!("{:?}", rect.left_top)
};
f(c); // Prints "Point { x: 2, y: 1 }".
```

generates a closure type roughly like the following:

// Note: This is not exactly how it is translated, this is only
for

// illustration.

```
struct Closure<'a> {
    left_top : &'a mut Point,
    right_bottom_x : &'a mut i32,
}
```

```
impl<'a> FnOnce<()> for Closure<'a> {
   type Output = String;
   extern "rust-call" fn call_once(self, args: ()) -> String {
      self.left_top.x += 1;
      *self.right_bottom_x += 1;
      format!("{:?}", self.left_top)
   }
}
so that the call to f works as if it were:
```

```
f(Closure{ left_top: &mut rect.left_top, right_bottom_x: &mut
rect.right_bottom.x });
```

```
r[type.closure.capture]
```

## **Capture modes**

r[type.closure.capture.intro] A *capture mode* determines how a <u>place</u> <u>expression</u> from the environment is borrowed or moved into the closure. The capture modes are:

- 1. Immutable borrow (ImmBorrow) --- The place expression is captured as a <u>shared reference</u>.
- 2. Unique immutable borrow (UniqueImmBorrow) --- This is similar to an immutable borrow, but must be unique as described <u>below</u>.
- 3. Mutable borrow (MutBorrow) --- The place expression is captured as a <u>mutable reference</u>.
- 4. Move (ByValue) --- The place expression is captured by <u>moving the</u> <u>value</u> into the closure.

r[type.closure.capture.precedence] Place expressions from the environment are captured from the first mode that is compatible with how the captured value is used inside the closure body. The mode is not affected by the code surrounding the closure, such as the lifetimes of involved variables or fields, or of the closure itself.

```
r[type.closure.capture.copy]
```

#### **Copy values**

Values that implement <u>Copy</u> that are moved into the closure are captured with the <u>ImmBorrow</u> mode.

```
let x = [0; 1024];
let c = || {
    let y = x; // x captured by ImmBorrow
};
```

r[type.closure.async.input]

## Async input capture

Async closures always capture all input arguments, regardless of whether or not they are used within the body.

# **Capture Precision**

r[type.closure.capture.precision.capture-path] A *capture path* is a sequence starting with a variable from the environment followed by zero or more place projections that were applied to that variable.

r[type.closure.capture.precision.place-projection] A *place projection* is a <u>field access</u>, <u>tuple index</u>, <u>dereference</u> (and automatic dereferences), or <u>array</u> <u>or slice index</u> expression applied to a variable.

r[type.closure.capture.precision.intro] The closure borrows or moves the capture path, which may be truncated based on the rules described below.

For example:

```
struct SomeStruct {
    f1: (i32, i32),
}
let s = SomeStruct { f1: (1, 2) };
let c = || {
    let x = s.f1.1; // s.f1.1 captured by ImmBorrow
};
c();
```

Here the capture path is the local variable s, followed by a field access .f1, and then a tuple index .1. This closure captures an immutable borrow of s.f1.1.

r[type.closure.capture.precision.shared-prefix]

# **Shared** prefix

In the case where a capture path and one of the ancestor's of that path are both captured by a closure, the ancestor path is captured with the highest capture mode among the two captures, CaptureMode = max(AncestorCaptureMode, DescendantCaptureMode), using the strict weak ordering:

```
ImmBorrow < UniqueImmBorrow < MutBorrow < ByValue</pre>
```

Note that this might need to be applied recursively.

```
// In this example, there are three different capture paths
with a shared ancestor:
# fn move_value<T>(_: T){}
let s = String::from("S");
let t = (s, String::from("T"));
let mut u = (t, String::from("U"));
let c = || {
    println!("{:?}", u); // u captured by ImmBorrow
    u.1.truncate(0); // u.0 captured by MutBorrow
    move_value(u.0.0); // u.0.0 captured by ByValue
};
c();
```

Overall this closure will capture u by ByValue.

r[type.closure.capture.precision.dereference-shared]

#### **Rightmost shared reference truncation**

The capture path is truncated at the rightmost dereference in the capture path if the dereference is applied to a shared reference.

This truncation is allowed because fields that are read through a shared reference will always be read via a shared reference or a copy. This helps reduce the size of the capture when the extra precision does not yield any benefit from a borrow checking perspective.

The reason it is the *rightmost* dereference is to help avoid a shorter lifetime than is necessary. Consider the following example:

```
struct Int(i32);
struct B<'a>(&'a i32);
struct MyStruct<'a> {
    a: &'static Int,
    b: B<'a>,
}
fn foo<'a, 'b>(m: &'a MyStruct<'b>) -> impl FnMut() + 'static
```

```
{
    let c = || drop(&m.a.0);
    c
}
```

If this were to capture m, then the closure would no longer outlive 'static, since m is constrained to 'a. Instead, it captures (\*(\*m).a) by ImmBorrow.

r[type.closure.capture.precision.wildcard]

# Wildcard pattern bindings

Closures only capture data that needs to be read. Binding a value with a <u>wildcard pattern</u> does not count as a read, and thus won't be captured. For example, the following closures will not capture  $\times$ :

```
let x = String::from("hello");
let c = || {
    let _ = x; // x is not captured
};
c();
let c = || match x { // x is not captured
    _ => println!("Hello World!")
};
c();
```

This also includes destructuring of tuples, structs, and enums. Fields matched with the [RestPattern] or [StructPatternEtCetera] are also not considered as read, and thus those fields will not be captured. The following illustrates some of these:

```
let x = (String::from("a"), String::from("b"));
let c = || {
    let (first, ..) = x; // captures `x.0` ByValue
};
// The first tuple field has been moved into the closure.
// The second tuple field is still accessible.
```

```
println!("{:?}", x.1);
c();
struct Example {
    f1: String,
    f2: String,
}
let e = Example {
    f1: String::from("first"),
    f2: String::from("second"),
};
let c = || {
    let Example { f2, .. } = e; // captures `e.f2` ByValue
};
// Field f2 cannot be accessed since it is moved into the
closure.
// Field f1 is still accessible.
println!("{:?}", e.f1);
c();
```

r[type.closure.capture.precision.wildcard.array-slice] Partial captures of arrays and slices are not supported; the entire slice or array is always captured even if used with wildcard pattern matching, indexing, or subslicing. For example:

```
#[derive(Debug)]
struct Example;
let x = [Example, Example];
let c = || {
    let [first, _] = x; // captures all of `x` ByValue
};
c();
println!("{:?}", x[1]); // ERROR: borrow of moved value: `x`
```

r[type.closure.capture.precision.wildcard.initialized] Values that are matched with wildcards must still be initialized.

```
let x: i32;
let c = || {
    let _ = x; // ERROR: used binding `x` isn't initialized
};
```

r[type.closure.capture.precision.move-dereference]

### **Capturing references in move contexts**

Because it is not allowed to move fields out of a reference, move closures will only capture the prefix of a capture path that runs up to, but not including, the first dereference of a reference. The reference itself will be moved into the closure.

```
struct T(String, String);
let mut t = T(String::from("foo"), String::from("bar"));
let t_mut_ref = &mut t;
let mut c = move || {
        t_mut_ref.0.push_str("123"); // captures `t_mut_ref`
ByValue
};
c();
```

r[type.closure.capture.precision.raw-pointer-dereference]

#### **Raw pointer dereference**

Because it is **unsafe** to dereference a raw pointer, closures will only capture the prefix of a capture path that runs up to, but not including, the first dereference of a raw pointer.

```
struct T(String, String);
let t = T(String::from("foo"), String::from("bar"));
let t_ptr = &t as *const T;
let c = || unsafe {
    println!("{}", (*t_ptr).0); // captures `t_ptr` by
ImmBorrow
```

}; c();

r[type.closure.capture.precision.union]

# **Union fields**

Because it is **unsafe** to access a union field, closures will only capture the prefix of a capture path that runs up to the union itself.

```
union U {
    a: (i32, i32),
    b: bool,
}
let u = U { a: (123, 456) };
let c = || {
    let x = unsafe { u.a.0 }; // captures `u` ByValue
};
c();
// This also includes writing to fields.
let mut u = U { a: (123, 456) };
let mut c = || {
    u.b = true; // captures `u` with MutBorrow
};
c();
```

r[type.closure.capture.precision.unaligned]

# **Reference into unaligned structs**

Because it is <u>undefined behavior</u> to create references to unaligned fields in a structure, closures will only capture the prefix of the capture path that runs up to, but not including, the first field access into a structure that uses <u>the packed representation</u>. This includes all fields, even those that are aligned, to protect against compatibility concerns should any of the fields in the structure change in the future.

```
#[repr(packed)]
struct T(i32, i32);
let t = T(2, 5);
let c = || {
    let a = t.0; // captures `t` with ImmBorrow
};
// Copies out of `t` are ok.
let (a, b) = (t.0, t.1);
c();
```

Similarly, taking the address of an unaligned field also captures the entire struct:

```
#[repr(packed)]
struct T(String, String);
let mut t = T(String::new(), String::new());
let c = || {
    let a = std::ptr::addr_of!(t.1); // captures `t` with
ImmBorrow
};
let a = t.0; // ERROR: cannot move out of `t.0` because it is
borrowed
c();
```

but the above works if it is not packed since it captures the field precisely:

```
struct T(String, String);
let mut t = T(String::new(), String::new());
let c = || {
    let a = std::ptr::addr_of!(t.1); // captures `t.1` with
ImmBorrow
};
// The move here is allowed.
```

```
let a = t.0;
c();
```

r[type.closure.capture.precision.box-deref]

#### Box vs other Deref implementations

The implementation of the <u>Deref</u> trait for <u>Box</u> is treated differently from other <u>Deref</u> implementations, as it is considered a special entity.

For example, let us look at examples involving Rc and Box. The \*rc is desugared to a call to the trait method deref defined on Rc, but since \*box is treated differently, it is possible to do a precise capture of the contents of the Box.

r[type.closure.capture.precision.box-non-move.not-moved]

#### Box with non-move closure

In a non-move closure, if the contents of the Box are not moved into the closure body, the contents of the Box are precisely captured.

```
struct S(String);
let b = Box::new(S(String::new()));
let c_box = || {
    let x = &(*b).0; // captures `(*b).0` by ImmBorrow
};
c_box();
// Contrast `Box` with another type that implements Deref:
let r = std::rc::Rc::new(S(String::new()));
let c_rc = || {
    let x = &(*r).0; // captures `r` by ImmBorrow
};
c_rc();
```

r[type.closure.capture.precision.box-non-move.moved] However, if the contents of the Box are moved into the closure, then the box is entirely

captured. This is done so the amount of data that needs to be moved into the closure is minimized.

```
// This is the same as the example above except the closure
// moves the value instead of taking a reference to it.
struct S(String);
let b = Box::new(S(String::new()));
let c_box = || {
    let x = (*b).0; // captures `b` with ByValue
};
c_box();
```

r[type.closure.capture.precision.box-move.read]

#### Box with move closure

Similarly to moving contents of a Box in a non-move closure, reading the contents of a Box in a move closure will capture the Box entirely.

```
struct S(i32);
let b = Box::new(S(10));
let c_box = move || {
    let x = (*b).0; // captures `b` with ByValue
};
```

r[type.closure.unique-immutable]

#### **Unique immutable borrows in captures**

Captures can occur by a special kind of borrow called a *unique immutable borrow*, which cannot be used anywhere else in the language and cannot be written out explicitly. It occurs when modifying the referent of a mutable reference, as in the following example:

```
let mut b = false;
let x = &mut b;
let mut c = || {
    // An ImmBorrow and a MutBorrow of `x`.
    let a = &x;
    *x = true; // `x` captured by UniqueImmBorrow
};
// The following line is an error:
// let y = &x;
c();
// However, the following is OK.
let z = &x;
```

In this case, borrowing × mutably is not possible, because × is not mut. But at the same time, borrowing × immutably would make the assignment illegal, because a & &mut reference might not be unique, so it cannot safely be used to modify a value. So a unique immutable borrow is used: it borrows × immutably, but like a mutable borrow, it must be unique.

In the above example, uncommenting the declaration of y will produce an error because it would violate the uniqueness of the closure's borrow of x; the declaration of z is valid because the closure's lifetime has expired at the end of the block, releasing the borrow.

r[type.closure.call]

## **Call traits and coercions**

r[type.closure.call.intro] Closure types all implement [FnOnce], indicating that they can be called once by consuming ownership of the closure. Additionally, some closures implement more specific call traits:

r[type.closure.call.fn-mut]

• A closure which does not move out of any captured variables implements [FnMut], indicating that it can be called by mutable reference.

r[type.closure.call.fn]

• A closure which does not mutate or move out of any captured variables implements [Fn], indicating that it can be called by shared reference.

[!NOTE] move closures may still implement [Fn] or [FnMut], even though they capture variables by move. This is because the traits implemented by a closure type are determined by what the closure does with captured values, not how it captures them.

r[type.closure.non-capturing] *Non-capturing closures* are closures that don't capture anything from their environment. Non-async, non-capturing closures can be coerced to function pointers (e.g., fn()) with the matching signature.

```
let add = |x, y| x + y;
let mut x = add(5,7);
type Binop = fn(i32, i32) -> i32;
let bo: Binop = add;
x = bo(5,7);
```

r[type.closure.async.traits]

## Async closure traits

r[type.closure.async.traits.fn-family] Async closures have a further restriction of whether or not they implement [FnMut] or [Fn].

The [Future] returned by the async closure has similar capturing characteristics as a closure. It captures place expressions from the async closure based on how they are used. The async closure is said to be *lending* to its [Future] if it has either of the following properties:

- The Future includes a mutable capture.
- The async closure captures by value, except when the value is accessed with a dereference projection.

If the async closure is lending to its Future, then [FnMut] and [Fn] are *not* implemented. [FnOnce] is always implemented.

**Example**: The first clause for a mutable capture can be illustrated with the following:

```
fn takes_callback<Fut: Future>(c: impl FnMut() -> Fut) {}
```

```
fn f() {
    let mut x = 1i32;
    let c = async || {
        x = 2; // x captured with MutBorrow
    };
        takes_callback(c); // ERROR: async closure does not
implement `FnMut`
}
```

The second clause for a regular value capture can be illustrated with the following:

```
fn takes_callback<Fut: Future>(c: impl Fn() -> Fut) {}
```

```
fn f() {
    let x = &1i32;
    let c = async move || {
        let a = x + 2; // x captured ByValue
    };
        takes_callback(c); // ERROR: async closure does not
```

implement `Fn`
}

The exception of the the second clause can be illustrated by using a dereference, which does allow Fn and FnMut to be implemented:

```
fn takes_callback<Fut: Future>(c: impl Fn() -> Fut) {}
fn f() {
    let x = &1i32;
    let c = async move || {
        let a = *x + 2;
    };
    takes_callback(c); // OK: implements `Fn`
}
```

r[type.closure.async.traits.async-family] Async closures implement [AsyncFn], [AsyncFnMut], and [AsyncFnOnce] in an analogous way as regular closures implement [Fn], [FnMut], and [FnOnce]; that is, depending on the use of the captured variables in its body.

r[type.closure.traits]

# **Other traits**

r[type.closure.traits.intro] All closure types implement <u>Sized</u>. Additionally, closure types implement the following traits if allowed to do so by the types of the captures it stores:

- <u>Clone</u>
- <u>Copy</u>
- <u>Sync</u>
- <u>Send</u>

r[type.closure.traits.behavior] The rules for <u>Send</u> and <u>Sync</u> match those for normal struct types, while <u>Clone</u> and <u>Copy</u> behave as if <u>derived</u>. For <u>Clone</u>, the order of cloning of the captured values is left unspecified.

Because captures are often by reference, the following general rules arise:

- A closure is <u>Sync</u> if all captured values are <u>Sync</u>.
- A closure is <u>Send</u> if all values captured by non-unique immutable reference are <u>Sync</u>, and all values captured by unique immutable or mutable reference, copy, or move are <u>Send</u>.
- A closure is <u>Clone</u> or <u>Copy</u> if it does not capture any values by unique immutable or mutable reference, and if all values it captures by copy or move are <u>Clone</u> or <u>Copy</u>, respectively.

r[type.closure.drop-order]

### **Drop Order**

If a closure captures a field of a composite types such as structs, tuples, and enums by value, the field's lifetime would now be tied to the closure. As a result, it is possible for disjoint fields of a composite types to be dropped at different times.

r[type.closure.capture.precision.edition2018.entirety]

# **Edition 2018 and before**

# **Closure types difference**

In Edition 2018 and before, closures always capture a variable in its entirety, without its precise capture path. This means that for the example used in the <u>Closure types</u> section, the generated closure type would instead look something like this:

```
struct Closure<'a> {
    rect : &'a mut Rectangle,
}
impl<'a> FnOnce<()> for Closure<'a> {
    type Output = String;
    extern "rust-call" fn call_once(self, args: ()) -> String {
        self.rect.left_top.x += 1;
        self.rect.right_bottom.x += 1;
        format!("{:?}", self.rect.left_top)
    }
}
and the call to f would work as follows:
```

```
f(Closure { rect: rect });
```

r[type.closure.capture.precision.edition2018.composite]

# **Capture precision difference**

Composite types such as structs, tuples, and enums are always captured in its entirety, not by individual fields. As a result, it may be necessary to borrow into a local variable in order to capture a single field:

```
# use std::collections::HashSet;
#
struct SetVec {
   set: HashSet<u32>,
   vec: Vec<u32>
}
```

```
impl SetVec {
    fn populate(&mut self) {
        let vec = &mut self.vec;
        self.set.iter().for_each(|&n| {
            vec.push(n);
        })
    }
}
```

If, instead, the closure were to use self.vec directly, then it would attempt to capture self by mutable reference. But since self.set is already borrowed to iterate over, the code would not compile.

r[type.closure.capture.precision.edition2018.move] If the move keyword is used, then all captures are by move or, for Copy types, by copy, regardless of whether a borrow would work. The move keyword is usually used to allow the closure to outlive the captured values, such as if the closure is being returned or used to spawn a new thread.

r[type.closure.capture.precision.edition2018.wildcard] Regardless of if the data will be read by the closure, i.e. in case of wild card patterns, if a variable defined outside the closure is mentioned within the closure the variable will be captured in its entirety.

r[type.closure.capture.precision.edition2018.drop-order]

## **Drop order difference**

As composite types are captured in their entirety, a closure which captures one of those composite types by value would drop the entire captured variable at the same time as the closure gets dropped.

```
{
    let tuple =
      (String::from("foo"), String::from("bar"));
    {
      let c = || { // ----+
      // tuple is captured into the closure |
      drop(tuple.0); // |
```

```
}; // |
} // 'c' and 'tuple' dropped here -----+
}
```

r[type.pointer]

# **Pointer types**

r[type.pointer.intro] All pointers are explicit first-class values. They can be moved or copied, stored into data structs, and returned from functions. r[type.pointer.reference]

## References (& and &mut)

r[type.pointer.reference.syntax]
ReferenceType -> `&` Lifetime? `mut`? TypeNoBounds

r[type.pointer.reference.shared]

#### Shared references (&)

r[type.pointer.reference.shared.intro] Shared references point to memory which is owned by some other value.

r[type.pointer.reference.shared.constraint-mutation] When a shared reference to a value is created, it prevents direct mutation of the value. <u>Interior mutability</u> provides an exception for this in certain circumstances. As the name suggests, any number of shared references to a value may exist. A shared reference type is written &type, or &'a type when you need to specify an explicit lifetime.

r[type.pointer.reference.shared.copy] Copying a reference is a "shallow" operation: it involves only copying the pointer itself, that is, pointers are Copy. Releasing a reference has no effect on the value it points to, but referencing of a <u>temporary value</u> will keep it alive during the scope of the reference itself.

r[type.pointer.reference.mut]

#### Mutable references (&mut)

r[type.pointer.reference.mut.intro] Mutable references point to memory which is owned by some other value. A mutable reference type is written &mut type or &'a mut type.

r[type.pointer.reference.mut.copy] A mutable reference (that hasn't been borrowed) is the only way to access the value it points to, so is not Copy.

r[type.pointer.raw]

#### Raw pointers (\*const and \*mut)

r[type.pointer.raw.syntax]
RawPointerType -> `\*` ( `mut` | `const` ) TypeNoBounds

r[type.pointer.raw.intro] Raw pointers are pointers without safety or liveness guarantees. Raw pointers are written as \*const T or \*mut T. For example \*const i32 means a raw pointer to a 32-bit integer.

r[type.pointer.raw.copy] Copying or dropping a raw pointer has no effect on the lifecycle of any other value.

r[type.pointer.raw.safety] Dereferencing a raw pointer is an <u>unsafe</u> <u>operation</u>.

This can also be used to convert a raw pointer to a reference by reborrowing it (&\* or &mut \*). Raw pointers are generally discouraged; they exist to support interoperability with foreign code, and writing performance-critical or low-level functions.

r[type.pointer.raw.cmp] When comparing raw pointers they are compared by their address, rather than by what they point to. When comparing raw pointers to <u>dynamically sized types</u> they also have their additional data compared.

r[type.pointer.raw.constructor] Raw pointers can be created directly using &raw const for \*const pointers and &raw mut for \*mut pointers.

r[type.pointer.smart]

#### **Smart Pointers**

The standard library contains additional 'smart pointer' types beyond references and raw pointers.

r[type.pointer.validity]

## **Bit validity**

r[type.pointer.validity.pointer-fragment] Despite pointers and references being similar to usize s in the machine code emitted on most platforms, the semantics of transmuting a reference or pointer type to a non-pointer type is currently undecided. Thus, it may not be valid to transmute a pointer or reference type, P, to a [u8; size\_of::<P>()].

r[type.pointer.validity.raw] For thin raw pointers (i.e., for P = \*const T or P = \*mut T for T: Sized), the inverse direction (transmuting from an integer or array of integers to P) is always valid. However, the pointer produced via such a transmutation may not be dereferenced (not even if T has size zero).

r[type.fn-pointer]

# **Function pointer types**

MaybeNamedFunctionParametersVariadic ->

( MaybeNamedParam `,` )\* MaybeNamedParam `,`
OuterAttribute\* `...`

r[type.fn-pointer.intro] Function pointer types, written using the fn keyword, refer to a function whose identity is not necessarily known at compile-time.

An example where Binop is defined as a function pointer type:

```
fn add(x: i32, y: i32) -> i32 {
    x + y
}
```
```
let mut x = add(5,7);
type Binop = fn(i32, i32) -> i32;
let bo: Binop = add;
x = bo(5,7);
```

r[type.fn-pointer.coercion] Function pointers can be created via a coercion from both <u>function items</u> and non-capturing, non-async <u>closures</u>.

r[type.fn-pointer.qualifiers] The unsafe qualifier indicates that the type's value is an <u>unsafe function</u>, and the extern qualifier indicates it is an <u>extern function</u>.

r[type.fn-pointer.constraint-variadic] Variadic parameters can only be specified with <u>extern</u> function types with the "C" or "cdecl" calling convention.

This also includes the corresponding [-unwind variants] [items.fn.extern.unwind].

r[type.fn-pointer.attributes]

### **Attributes on function pointer parameters**

Attributes on function pointer parameters follow the same rules and restrictions as <u>regular function parameters</u>.

r[type.trait-object]

## **Trait objects**

```
r[type.trait-object.syntax]
TraitObjectType -> `dyn`? TypeParamBounds
```

```
TraitObjectTypeOneBound -> `dyn`? TraitBound
```

r[type.trait-object.intro] A *trait object* is an opaque value of another type that implements a set of traits. The set of traits is made up of a <u>dyn</u> <u>compatible</u> *base trait* plus any number of <u>auto traits</u>.

r[type.trait-object.impls] Trait objects implement the base trait, its auto traits, and any <u>supertraits</u> of the base trait.

r[type.trait-object.name] Trait objects are written as the keyword dyn followed by a set of trait bounds, but with the following restrictions on the trait bounds.

r[type.trait-object.constraint] There may not be more than one non-auto trait, no more than one lifetime, and opt-out bounds (e.g. ?Sized) are not allowed. Furthermore, paths to traits may be parenthesized.

For example, given a trait Trait, the following are all trait objects:

```
• dyn Trait
```

- dyn Trait + Send
- dyn Trait + Send + Sync
- dyn Trait + 'static
- dyn Trait + Send + 'static
- dyn Trait +
- dyn 'static + Trait.
- dyn (Trait)

r[type.trait-object.syntax-edition2021]

[!EDITION-2021] Before the 2021 edition, the dyn keyword may be omitted.

r[type.trait-object.syntax-edition2018]

[!EDITION-2018] In the 2015 edition, if the first bound of the trait object is a path that starts with ::, then the dyn will be treated as a part of the path. The first path can be put in parenthesis to get around this. As such, if you want a trait object with the trait ::your\_module::Trait, you should write it as dyn (::your\_module::Trait).

Beginning in the 2018 edition, dyn is a true keyword and is not allowed in paths, so the parentheses are not necessary.

r[type.trait-object.alias] Two trait object types alias each other if the base traits alias each other and if the sets of auto traits are the same and the lifetime bounds are the same. For example, dyn Trait + Send + UnwindSafe is the same as dyn Trait + UnwindSafe + Send.

r[type.trait-object.unsized] Due to the opaqueness of which concrete type the value is of, trait objects are <u>dynamically sized types</u>. Like all <u>DSTs</u>, trait objects are used behind some type of pointer; for example &dyn SomeTrait or Box<dyn SomeTrait>. Each instance of a pointer to a trait object includes:

- a pointer to an instance of a type T that implements SomeTrait
- a virtual method table, often just called a vtable, which contains, for each method of SomeTrait and its <u>supertraits</u> that T implements, a pointer to T's implementation (i.e. a function pointer).

The purpose of trait objects is to permit "late binding" of methods. Calling a method on a trait object results in virtual dispatch at runtime: that is, a function pointer is loaded from the trait object vtable and invoked indirectly. The actual implementation for each vtable entry can vary on an object-by-object basis.

An example of a trait object:

```
trait Printable {
    fn stringify(&self) -> String;
}
impl Printable for i32 {
```

```
fn stringify(&self) -> String { self.to_string() }
fn print(a: Box<dyn Printable>) {
    println!("{}", a.stringify());
}
fn main() {
    print(Box::new(10) as Box<dyn Printable>);
}
```

In this example, the trait **Printable** occurs as a trait object in both the type signature of print, and the cast expression in main.

r[type.trait-object.lifetime-bounds]

### **Trait Object Lifetime Bounds**

Since a trait object can contain references, the lifetimes of those references need to be expressed as part of the trait object. This lifetime is written as Trait + 'a. There are <u>defaults</u> that allow this lifetime to usually be inferred with a sensible choice.

r[type.impl-trait]

## **Impl trait**

```
r[type.impl-trait.syntax]
ImplTraitType -> `impl` TypeParamBounds
```

```
ImplTraitTypeOneBound -> `impl` TraitBound
```

r[type.impl-trait.intro] impl Trait provides ways to specify unnamed but concrete types that implement a specific trait. It can appear in two sorts of places: argument position (where it can act as an anonymous type parameter to functions), and return position (where it can act as an abstract return type).

```
trait Trait {}
# impl Trait for () {}
// argument position: anonymous type parameter
fn foo(arg: impl Trait) {
}
// return position: abstract return type
fn bar() -> impl Trait {
}
```

r[type.impl-trait.param]

#### **Anonymous type parameters**

[!NOTE] This is often called "impl Trait in argument position". (The term "parameter" is more correct here, but "impl Trait in argument position" is the phrasing used during the development of this feature, and it remains in parts of the implementation.)

r[type.impl-trait.param.intro] Functions can use impl followed by a set of trait bounds to declare a parameter as having an anonymous type. The caller must provide a type that satisfies the bounds declared by the anonymous type parameter, and the function can only use the methods available through the trait bounds of the anonymous type parameter.

For example, these two forms are almost equivalent:

```
trait Trait {}
// generic type parameter
fn with_generic_type<T: Trait>(arg: T) {
}
// impl Trait in argument position
fn with_impl_trait(arg: impl Trait) {
}
```

r[type.impl-trait.param.generic] That is, impl Trait in argument position is syntactic sugar for a generic type parameter like <T: Trait>, except that the type is anonymous and doesn't appear in the [GenericParams] list.

[!NOTE] For function parameters, generic type parameters and impl Trait are not exactly equivalent. With a generic parameter such as <T: Trait>, the caller has the option to explicitly specify the generic argument for T at the call site using [GenericArgs], for example, foo::<usize>(1). Changing a parameter from either one to the other can constitute a breaking change for the callers of a function, since this changes the number of generic arguments.

r[type.impl-trait.return]

### Abstract return types

[!NOTE] This is often called "impl Trait in return position".

r[type.impl-trait.return.intro] Functions can use impl Trait to return an abstract return type. These types stand in for another concrete type where the caller may only use the methods declared by the specified Trait.

r[type.impl-trait.return.constraint-body] Each possible return value from the function must resolve to the same concrete type.

impl Trait in return position allows a function to return an unboxed abstract type. This is particularly useful with <u>closures</u> and iterators. For example, closures have a unique, un-writable type. Previously, the only way to return a closure from a function was to use a <u>trait object</u>:

```
fn returns_closure() -> Box<dyn Fn(i32) -> i32> {
    Box::new(|x| x + 1)
}
```

This could incur performance penalties from heap allocation and dynamic dispatch. It wasn't possible to fully specify the type of the closure, only to use the Fn trait. That means that the trait object is necessary. However, with impl Trait, it is possible to write this more simply:

which also avoids the drawbacks of using a boxed trait object.

Similarly, the concrete types of iterators could become very complex, incorporating the types of all previous iterators in a chain. Returning impl Iterator means that a function only exposes the Iterator trait as a bound on its return type, instead of explicitly specifying all of the other iterator types involved.

r[type.impl-trait.return-in-trait]

# Return-position impl Trait in traits and trait implementations

r[type.impl-trait.return-in-trait.intro] Functions in traits may also use impl Trait as a syntax for an anonymous associated type.

r[type.impl-trait.return-in-trait.desugaring] Every impl Trait in the return type of an associated function in a trait is desugared to an anonymous associated type. The return type that appears in the implementation's function signature is used to determine the value of the associated type.

r[type.impl-trait.generic-captures]

### Capturing

Behind each return-position impl Trait abstract type is some hidden concrete type. For this concrete type to use a generic parameter, that generic parameter must be *captured* by the abstract type.

r[type.impl-trait.generic-capture.auto]

### **Automatic capturing**

r[type.impl-trait.generic-capture.auto.intro] Return-position impl Trait abstract types automatically capture all in-scope generic parameters, including generic type, const, and lifetime parameters (including higherranked ones).

r[type.impl-trait.generic-capture.edition2024]

[!EDITION-2024] Before the 2024 edition, on free functions and on associated functions and methods of inherent impls, generic lifetime parameters that do not appear in the bounds of the abstract return type are not automatically captured.

r[type.impl-trait.generic-capture.precise]

### **Precise capturing**

r[type.impl-trait.generic-capture.precise.use] The set of generic parameters captured by a return-position impl Trait abstract type may be explicitly controlled with a <u>use<...> bound</u>. If present, only the generic parameters listed in the use<...> bound will be captured. E.g.:

r[type.impl-trait.generic-capture.precise.constraint-single] Currently, only one use<..> bound may be present in a bounds list, all in-scope type and const generic parameters must be included, and all lifetime parameters that appear in other bounds of the abstract type must be included.

r[type.impl-trait.generic-capture.precise.constraint-lifetime] Within the use<..> bound, any lifetime parameters present must appear before all type and const generic parameters, and the elided lifetime ('\_) may be present if it is otherwise allowed to appear within the impl Trait return type.

r[type.impl-trait.generic-capture.precise.constraint-param-impl-trait]
Because all in-scope type parameters must be included by name, a use<..>
bound may not be used in the signature of items that use argument-position
impl Trait, as those items have anonymous type parameters in scope.

r[type.impl-trait.generic-capture.precise.constraint-in-trait] Any
use<..> bound that is present in an associated function in a trait definition
must include all generic parameters of the trait, including the implicit Self
generic type parameter of the trait.

# Differences between generics and impl Trait in return position

In argument position, impl Trait is very similar in semantics to a generic type parameter. However, there are significant differences between the two in return position. With impl Trait, unlike with a generic type parameter, the function chooses the return type, and the caller cannot choose the return type.

The function:

```
# trait Trait {}
fn foo<T: Trait>() -> T {
    // ...
# panic!()
}
```

allows the caller to determine the return type,  $\top$ , and the function returns that type.

The function:

doesn't allow the caller to determine the return type. Instead, the function chooses the return type, but only promises that it will implement Trait.

```
r[type.impl-trait.constraint]
```

### Limitations

impl Trait can only appear as a parameter or return type of a nonextern function. It cannot be the type of a let binding, field type, or appear inside a type alias. r[type.generic]

### **Type parameters**

Within the body of an item that has type parameter declarations, the names of its type parameters are types:

```
fn to_vec<A: Clone>(xs: &[A]) -> Vec<A> {
    if xs.is_empty() {
        return vec![];
    }
    let first: A = xs[0].clone();
    let mut rest: Vec<A> = to_vec(&xs[1..]);
    rest.insert(0, first);
    rest
}
```

Here, first has type A, referring to to\_vec's A type parameter; and rest has type Vec<A>, a vector with element type A.

r[type.inferred]

## **Inferred type**

r[type.inferred.syntax]

InferredType -> `\_`

r[type.inferred.intro] The inferred type asks the compiler to infer the type if possible based on the surrounding information available.

r[type.inferred.constraint] It cannot be used in item signatures.

It is often used in generic arguments:

let x: Vec<\_> = (0..10).collect();

r[dynamic-sized]

## **Dynamically Sized Types**

r[dynamic-sized.intro] Most types have a fixed size that is known at compile time and implement the trait <u>Sized</u>. A type with a size that is known only at run-time is called a *dynamically sized type* (*DST*) or, informally, an unsized type. <u>Slices</u> and <u>trait objects</u> are two examples of <u>DSTs</u>.

r[dynamic-sized.restriction] Such types can only be used in certain cases:

r[dynamic-sized.pointer-types]

- <u>Pointer types</u> to <u>DSTs</u> are sized but have twice the size of pointers to sized types
  - Pointers to slices also store the number of elements of the slice.
  - Pointers to trait objects also store a pointer to a vtable.

r[dynamic-sized.question-sized]

• <u>DSTs</u> can be provided as type arguments to generic type parameters having the special <code>?Sized</code> bound. They can also be used for associated type definitions when the corresponding associated type declaration has a <code>?Sized</code> bound. By default, any type parameter or associated type has a <code>Sized</code> bound, unless it is relaxed using <code>?Sized</code>.

r[dynamic-sized.trait-impl]

• Traits may be implemented for <u>DSTs</u>. Unlike with generic type parameters, Self: ?Sized is the default in trait definitions.

r[dynamic-sized.struct-field]

• Structs may contain a <u>DST</u> as the last field; this makes the struct itself a <u>DST</u>.

[!NOTE] <u>Variables</u>, function parameters, <u>const</u> items, and <u>static</u> items must be <u>Sized</u>.

r[layout]

### **Type Layout**

r[layout.intro] The layout of a type is its size, alignment, and the relative offsets of its fields. For enums, how the discriminant is laid out and interpreted is also part of type layout.

r[layout.guarantees] Type layout can be changed with each compilation. Instead of trying to document exactly what is done, we only document what is guaranteed today.

Note that even types with the same layout can still differ in how they are passed across function boundaries. For function call ABI compatibility of types, see <u>here</u>.

r[layout.properties]

#### **Size and Alignment**

All values have an alignment and size.

r[layout.properties.align] The *alignment* of a value specifies what addresses are valid to store the value at. A value of alignment **n** must only be stored at an address that is a multiple of n. For example, a value with an alignment of 2 must be stored at an even address, while a value with an alignment of 1 can be stored at any address. Alignment is measured in bytes, and must be at least 1, and always a power of 2. The alignment of a value can be checked with the <u>align of val</u> function.

r[layout.properties.size] The *size* of a value is the offset in bytes between successive elements in an array with that item type including alignment padding. The size of a value is always a multiple of its alignment. Note that some types are zero-sized; 0 is considered a multiple of any alignment (for example, on some platforms, the type [u16; 0] has size 0 and alignment 2). The size of a value can be checked with the <u>size of val</u> function.

r[layout.properties.sized] Types where all values have the same size and alignment, and both are known at compile time, implement the <u>Sized</u> trait and can be checked with the <u>size of</u> and <u>align of</u> functions. Types that are not <u>Sized</u> are known as <u>dynamically sized types</u>. Since all values of a <u>Sized</u> type share the same size and alignment, we refer to those shared values as the size of the type and the alignment of the type respectively.

r[layout.primitive]

### **Primitive Data Layout**

Туре	<pre>size_of::<type>()</type></pre>
bool	1
u8 / i8	1
u16 / i16	2
u32 / i32	4
u64 / i64	8
u128 / i128	16
usize / isize	See below
f32	4
f64	8
char	4

r[layout.primitive.size] The size of most primitives is given in this table.

r[layout.primitive.size-int] usize and isize have a size big enough to contain every address on the target platform. For example, on a 32 bit target, this is 4 bytes, and on a 64 bit target, this is 8 bytes.

r[layout.primitive.align] The alignment of primitives is platformspecific. In most cases, their alignment is equal to their size, but it may be less. In particular, i128 and u128 are often aligned to 4 or 8 bytes even though their size is 16, and on many 32-bit platforms, i64, u64, and f64 are only aligned to 4 bytes, not 8.

r[layout.pointer]

### **Pointers and References Layout**

r[layout.pointer.intro] Pointers and references have the same layout. Mutability of the pointer or reference does not change the layout.

r[layout.pointer.thin] Pointers to sized types have the same size and alignment as usize.

r[layout.pointer.unsized] Pointers to unsized types are sized. The size and alignment is guaranteed to be at least equal to the size and alignment of a pointer.

[!NOTE] Though you should not rely on this, all pointers to <u>DSTs</u> are currently twice the size of the size of **usize** and have the same alignment.

r[layout.array]

#### **Array Layout**

An array of [T; N] has a size of  $size_of::<T>() * N$  and the same alignment of T. Arrays are laid out so that the zero-based nth element of the array is offset from the start of the array by  $n * size_of::<T>()$  bytes.

r[layout.slice]

### **Slice Layout**

Slices have the same layout as the section of the array they slice.

[!NOTE] This is about the raw [T] type, not pointers (&[T], Box<[T]>, etc.) to slices.

r[layout.str]

### str Layout

String slices are a UTF-8 representation of characters that have the same layout as slices of type [u8].

r[layout.tuple]

#### **Tuple Layout**

r[layout.tuple.general] Tuples are laid out according to the <u>Rust</u> <u>representation</u>.

r[layout.tuple.unit] The exception to this is the unit tuple (()), which is guaranteed as a zero-sized type to have a size of 0 and an alignment of 1.

r[layout.trait-object]

### **Trait Object Layout**

Trait objects have the same layout as the value the trait object is of.

[!NOTE] This is about the raw trait object types, not pointers (&dyn Trait, Box<dyn Trait>, etc.) to trait objects.

r[layout.closure]

### **Closure Layout**

Closures have no layout guarantees. r[layout.repr]

### Representations

r[layout.repr.intro] All user-defined composite types (structs, enums, and unions) have a *representation* that specifies what the layout is for the type.

r[layout.repr.kinds] The possible representations for a type are:

- <u>Rust</u> (default)
- <u>C</u>
- The <u>primitive representations</u>
- <u>transparent</u>

r[layout.repr.attribute] The representation of a type can be changed by applying the **repr** attribute to it. The following example shows a struct with a **C** representation.

```
#[repr(C)]
struct ThreeInts {
    first: i16,
    second: i8,
    third: i32
}
```

r[layout.repr.align-packed] The alignment may be raised or lowered with the align and packed modifiers respectively. They alter the representation specified in the attribute. If no representation is specified, the default one is altered.

```
// Default representation, alignment lowered to 2.
#[repr(packed(2))]
struct PackedStruct {
   first: i16,
   second: i8,
   third: i32
}
// C representation, alignment raised to 8
```
```
#[repr(C, align(8))]
struct AlignedStruct {
    first: i16,
    second: i8,
    third: i32
}
```

}

[!NOTE] As a consequence of the representation being an attribute on the item, the representation does not depend on generic parameters. Any two types with the same name have the same representation. For example, Foo<Bar> and Foo<Baz> both have the same representation.

r[layout.repr.inter-field] The representation of a type can change the padding between fields, but does not change the layout of the fields themselves. For example, a struct with a C representation that contains a struct Inner with the Rust representation will not change the layout of Inner.

r[layout.repr.rust]

## The Rust Representation

r[layout.repr.rust.intro] The Rust representation is the default representation for nominal types without a repr attribute. Using this representation explicitly through a repr attribute is guaranteed to be the same as omitting the attribute entirely.

r[layout.repr.rust.layout] The only data layout guarantees made by this representation are those required for soundness. They are:

- 1. The fields are properly aligned.
- 2. The fields do not overlap.
- 3. The alignment of the type is at least the maximum alignment of its fields.

r[layout.repr.rust.alignment] Formally, the first guarantee means that the offset of any field is divisible by that field's alignment.

r[layout.repr.rust.field-storage] The second guarantee means that the fields can be ordered such that the offset plus the size of any field is less

than or equal to the offset of the next field in the ordering. The ordering does not have to be the same as the order in which the fields are specified in the declaration of the type.

Be aware that the second guarantee does not imply that the fields have distinct addresses: zero-sized types may have the same address as other fields in the same struct.

r[layout.repr.rust.unspecified] There are no other guarantees of data layout made by this representation.

r[layout.repr.c]

### The c Representation

r[layout.repr.c.intro] The **c** representation is designed for dual purposes. One purpose is for creating types that are interoperable with the C Language. The second purpose is to create types that you can soundly perform operations on that rely on data layout such as reinterpreting values as a different type.

Because of this dual purpose, it is possible to create types that are not useful for interfacing with the C programming language.

r[layout.repr.c.constraint] This representation can be applied to structs, unions, and enums. The exception is <u>zero-variant enums</u> for which the <u>c</u> representation is an error.

r[layout.repr.c.struct]

#### #[repr(C)] Structs

r[layout.repr.c.struct.align] The alignment of the struct is the alignment of the most-aligned field in it.

r[layout.repr.c.struct.size-field-offset] The size and offset of fields is determined by the following algorithm.

Start with a current offset of 0 bytes.

For each field in declaration order in the struct, first determine the size and alignment of the field. If the current offset is not a multiple of the field's alignment, then add padding bytes to the current offset until it is a multiple of the field's alignment. The offset for the field is what the current offset is now. Then increase the current offset by the size of the field. Finally, the size of the struct is the current offset rounded up to the nearest multiple of the struct's alignment.

```
Here is this algorithm described in pseudocode.
/// Returns the amount of padding needed after `offset` to
ensure that the
/// following address will be aligned to `alignment`.
fn padding_needed_for(offset: usize, alignment: usize) -> usize
{
    let misalignment = offset % alignment;
    if misalignment > 0 {
        // round up to next multiple of `alignment`
        alignment - misalignment
    } else {
        // already a multiple of `alignment`
        0
    }
}
struct.alignment
                                     struct.fields().map(|field|
                          =
field.alignment).max();
let current offset = 0;
for field in struct.fields_in_declaration_order() {
    // Increase the current offset so that it's a multiple of
the alignment
    // of this field. For the first field, this will always be
zero.
    // The skipped bytes are called padding bytes.
         current_offset += padding_needed_for(current_offset,
field.alignment);
    struct[field].offset = current_offset;
    current_offset += field.size;
```

}

```
struct.size = current_offset
padding_needed_for(current_offset, struct.alignment);
```

[!WARNING] This pseudocode uses a naive algorithm that ignores overflow issues for the sake of clarity. To perform memory layout computations in actual code, use <u>Layout</u>.

[!NOTE] This algorithm can produce zero-sized structs. In C, an empty struct declaration like struct Foo { } is illegal. However, both gcc and clang support options to enable such structs, and assign them size zero. C++, in contrast, gives empty structs a size of 1, unless they are inherited from or they are fields that have the [[no\_unique\_address]] attribute, in which case they do not increase the overall size of the struct.

r[layout.repr.c.union]

#### #[repr(C)] Unions

r[layout.repr.c.union.intro] A union declared with #[repr(C)] will have the same size and alignment as an equivalent C union declaration in the C language for the target platform.

r[layout.repr.c.union.size-align] The union will have a size of the maximum size of all of its fields rounded to its alignment, and an alignment of the maximum alignment of all of its fields. These maximums may come from different fields.

```
#[repr(C)]
union Union {
    f1: u16,
    f2: [u8; 4],
}
assert_eq!(std::mem::size_of::<Union>(), 4); // From f2
assert_eq!(std::mem::align_of::<Union>(), 2); // From f1
#[repr(C)]
```

+

r[layout.repr.c.enum]

### #[repr(C)] Field-less Enums

For <u>field-less enums</u>, the <u>C</u> representation has the size and alignment of the default <u>enum</u> size and alignment for the target platform's C ABI.

[!NOTE] The enum representation in C is implementation defined, so this is really a "best guess". In particular, this may be incorrect when the C code of interest is compiled with certain flags.

[!WARNING] There are crucial differences between an enum in the C language and Rust's <u>field-less enums</u> with this representation. An enum in C is mostly a typedef plus some named constants; in other words, an object of an enum type can hold any integer value. For example, this is often used for bitflags in C. In contrast, Rust's <u>field-less enums</u> can only legally hold the discriminant values, everything else is <u>undefined behavior</u>. Therefore, using a field-less enum in FFI to model a C enum is often wrong.

r[layout.repr.c.adt]

### #[repr(C)] Enums With Fields

r[layout.repr.c.adt.intro] The representation of a repr(C) enum with fields is a repr(C) struct with two fields, also called a "tagged union" in C:

r[layout.repr.c.adt.tag]

- a repr(C) version of the enum with all fields removed ("the tag") r[layout.repr.c.adt.fields]
- a repr(C) union of repr(C) structs for the fields of each variant that had them ("the payload")

[!NOTE] Due to the representation of repr(C) structs and unions, if a variant has a single field there is no difference between putting that field directly in the union or wrapping it in a struct; any system which wishes to manipulate such an enum's representation may therefore use whichever form is more convenient or consistent for them.

```
// This Enum has the same representation as ...
#[repr(C)]
enum MyEnum {
    A(u32),
    B(f32, u64),
    C { x: u32, y: u8 },
    D,
 }
// ... this struct.
#[repr(C)]
struct MyEnumRepr {
    taq: MyEnumDiscriminant,
    payload: MyEnumFields,
}
// This is the discriminant enum.
#[repr(C)]
enum MyEnumDiscriminant { A, B, C, D }
```

```
// This is the variant union.
#[repr(C)]
union MyEnumFields {
    A: MyAFields,
    B: MyBFields,
    C: MyCFields,
    D: MyDFields,
}
#[repr(C)]
#[derive(Copy, Clone)]
struct MyAFields(u32);
#[repr(C)]
#[derive(Copy, Clone)]
struct MyBFields(f32, u64);
#[repr(C)]
#[derive(Copy, Clone)]
struct MyCFields { x: u32, y: u8 }
// This struct could be omitted (it is a zero-sized type), and
it must be in
// C/C++ headers.
#[repr(C)]
#[derive(Copy, Clone)]
struct MyDFields;
```

r[layout.repr.primitive]

### **Primitive representations**

r[layout.repr.primitive.intro] The *primitive representations* are the representations with the same names as the primitive integer types. That is: u8, u16, u32, u64, u128, usize, i8, i16, i32, i64, i128, and isize.

r[layout.repr.primitive.constraint] Primitive representations can only be applied to enumerations and have different behavior whether the enum has fields or no fields. It is an error for <u>zero-variant enums</u> to have a primitive representation. Combining two primitive representations together is an error.

r[layout.repr.primitive.enum]

### **Primitive Representation of Field-less Enums**

For <u>field-less enums</u>, primitive representations set the size and alignment to be the same as the primitive type of the same name. For example, a fieldless enum with a us representation can only have discriminants between 0 and 255 inclusive.

r[layout.repr.primitive.adt]

### **Primitive Representation of Enums With Fields**

The representation of a primitive representation enum is a repr(C) union of repr(C) structs for each variant with a field. The first field of each struct in the union is the primitive representation version of the enum with all fields removed ("the tag") and the remaining fields are the fields of that variant.

[!NOTE] This representation is unchanged if the tag is given its own member in the union, should that make manipulation more clear for you (although to follow the C++ standard the tag member should be wrapped in a struct).

```
// This enum has the same representation as ...
#[repr(u8)]
enum MyEnum {
        A(u32),
        B(f32, u64),
        C { x: u32, y: u8 },
        D,
     }
// ... this union.
#[repr(C)]
union MyEnumRepr {
```

```
A: MyVariantA,
    B: MyVariantB,
    C: MyVariantC,
    D: MyVariantD,
}
// This is the discriminant enum.
#[repr(u8)]
#[derive(Copy, Clone)]
enum MyEnumDiscriminant { A, B, C, D }
#[repr(C)]
#[derive(Clone, Copy)]
struct MyVariantA(MyEnumDiscriminant, u32);
#[repr(C)]
#[derive(Clone, Copy)]
struct MyVariantB(MyEnumDiscriminant, f32, u64);
#[repr(C)]
#[derive(Clone, Copy)]
struct MyVariantC { tag: MyEnumDiscriminant, x: u32, y: u8 }
#[repr(C)]
#[derive(Clone, Copy)]
struct MyVariantD(MyEnumDiscriminant);
```

r[layout.repr.primitive-c]

# Combining primitive representations of enums with fields and #[repr(C)]

For enums with fields, it is also possible to combine repr(C) and a primitive representation (e.g., repr(C, u8)). This modifies the repr(C) by changing the representation of the discriminant enum to the chosen primitive instead. So, if you chose the u8 representation, then the discriminant enum would have a size and alignment of 1 byte.

The discriminant enum from the example <u>earlier</u> then becomes:

```
#[repr(C, u8)] // `u8` was added
enum MyEnum {
        A(u32),
        B(f32, u64),
        C { x: u32, y: u8 },
        D,
     }
// ...
#[repr(u8)] // So `u8` is used here instead of `C`
enum MyEnumDiscriminant { A, B, C, D }
// ...
```

For example, with a repr(C, u8) enum it is not possible to have 257 unique discriminants ("tags") whereas the same enum with only a repr(C) attribute will compile without any problems.

Using a primitive representation in addition to repr(C) can change the size of an enum from the repr(C) form:

```
#[repr(C)]
enum EnumC {
    Variant0(u8),
    Variant1,
}
#[repr(C, u8)]
enum Enum8 {
    Variant0(u8),
    Variant1,
}
#[repr(C, u16)]
enum Enum16 {
```

```
Variant0(u8),
Variant1,
}
// The size of the C representation is platform dependent
assert_eq!(std::mem::size_of::<EnumC>(), 8);
// One byte for the discriminant and one byte for the value in
Enum8::Variant0
assert_eq!(std::mem::size_of::<Enum8>(), 2);
// Two bytes for the discriminant and one byte for the value
in Enum16::Variant0
// plus one byte of padding.
assert_eq!(std::mem::size_of::<Enum16>(), 4);
```

r[layout.repr.alignment]

## The alignment modifiers

r[layout.repr.alignment.intro] The align and packed modifiers can be used to respectively raise or lower the alignment of structs and unions. packed may also alter the padding between fields (although it will not alter the padding inside of any field). On their own, align and packed do not provide guarantees about the order of fields in the layout of a struct or the layout of an enum variant, although they may be combined with representations (such as C) which do provide such guarantees.

r[layout.repr.alignment.constraint-alignment] The alignment is specified as an integer parameter in the form of #[repr(align(x))] or #[repr(packed(x))]. The alignment value must be a power of two from 1 up to 2<sup>29</sup>. For packed, if no value is given, as in #[repr(packed)], then the value is 1.

r[layout.repr.alignment.align] For align, if the specified alignment is less than the alignment of the type without the align modifier, then the alignment is unaffected.

r[layout.repr.alignment.packed] For packed, if the specified alignment is greater than the type's alignment without the packed modifier, then the

alignment and layout is unaffected.

r[layout.repr.alignment.packed-fields] The alignments of each field, for the purpose of positioning fields, is the smaller of the specified alignment and the alignment of the field's type.

r[layout.repr.alignment.packed-padding] Inter-field padding is guaranteed to be the minimum required in order to satisfy each field's (possibly altered) alignment (although note that, on its own, packed does not provide any guarantee about field ordering). An important consequence of these rules is that a type with #[repr(packed(1))] (or # [repr(packed)]) will have no inter-field padding.

r[layout.repr.alignment.constraint-exclusive] The align and packed modifiers cannot be applied on the same type and a packed type cannot transitively contain another align ed type. align and packed may only be applied to the <u>Rust</u> and <u>C</u> representations.

r[layout.repr.alignment.enum] The align modifier can also be applied on an enum. When it is, the effect on the enum's alignment is the same as if the enum was wrapped in a newtype struct with the same align modifier.

[!NOTE] References to unaligned fields are not allowed because it is <u>undefined behavior</u>. When fields are unaligned due to an alignment modifier, consider the following options for using references and dereferences:

```
#[repr(packed)]
struct Packed {
    f1: u8,
    f2: u16,
}
let mut e = Packed { f1: 1, f2: 2 };
// Instead of creating a reference to a field, copy the
value to a local variable.
let x = e.f2;
// Or in situations like `println!` which creates a
reference, use braces
```

```
// to change it to a copy of the value.
println!("{}", {e.f2});
// Or if you need a pointer, use the unaligned methods
for reading and writing
// instead of dereferencing the pointer directly.
let ptr: *const u16 = &raw const e.f2;
let value = unsafe { ptr.read_unaligned() };
let mut_ptr: *mut u16 = &raw mut e.f2;
unsafe { mut_ptr.write_unaligned(3) }
```

r[layout.repr.transparent]

## The transparent Representation

r[layout.repr.transparent.constraint-field] The transparent representation can only be used on a struct or an enum with a single variant that has:

- any number of fields with size 0 and alignment 1 (e.g. <u>PhantomData<T></u>), and
- at most one other field.

r[layout.repr.transparent.layout-abi] Structs and enums with this representation have the same layout and ABI as the only non-size 0 non-alignment 1 field, if present, or unit otherwise.

This is different than the C representation because a struct with the C representation will always have the ABI of a C struct while, for example, a struct with the transparent representation with a primitive field will have the ABI of the primitive field.

r[layout.repr.transparent.constraint-exclusive] Because this representation delegates type layout to another type, it cannot be used with any other representation.

r[interior-mut]

## **Interior Mutability**

r[interior-mut.intro] Sometimes a type needs to be mutated while having multiple aliases. In Rust this is achieved using a pattern called *interior mutability*.

r[interior-mut.shared-ref] A type has interior mutability if its internal state can be changed through a <u>shared reference</u> to it.

r[interior-mut.no-constraint] This goes against the usual <u>requirement</u> that the value pointed to by a shared reference is not mutated.

r[interior-mut.unsafe-cell] [std::cell::UnsafeCell<T>] type is the only allowed way to disable this requirement. When UnsafeCell<T> is immutably aliased, it is still safe to mutate, or obtain a mutable reference to, the T it contains.

r[interior-mut.mut-unsafe-cell] As with all other types, it is undefined behavior to have multiple &mut UnsafeCell<T> aliases.

r[interior-mut.abstraction] Other types with interior mutability can be created by using UnsafeCell<T> as a field. The standard library provides a variety of types that provide safe interior mutability APIs.

r[interior-mut.ref-cell] For example, [std::cell::RefCell<T>] uses run-time borrow checks to ensure the usual rules around multiple references.

r[interior-mut.atomic] The [std::sync::atomic] module contains types that wrap a value that is only accessed with atomic operations, allowing the value to be shared and mutated across threads.

r[subtype]

## **Subtyping and Variance**

r[subtype.intro] Subtyping is implicit and can occur at any stage in type checking or inference.

r[subtype.kinds] Subtyping is restricted to two cases: variance with respect to lifetimes and between types with higher ranked lifetimes. If we were to erase lifetimes from types, then the only subtyping would be due to type equality.

Consider the following example: string literals always have 'static lifetime. Nevertheless, we can assign s to t:

```
fn bar<'a>() {
    let s: &'static str = "hi";
    let t: &'a str = s;
}
```

Since 'static outlives the lifetime parameter 'a, &'static str is a subtype of &'a str.

r[subtype.higher-ranked] <u>Higher-ranked</u> <u>function pointers</u> and <u>trait</u> <u>objects</u> have another subtype relation. They are subtypes of types that are given by substitutions of the higher-ranked lifetimes. Some examples:

```
// Here 'a is substituted for 'static
let subtype: &(for<'a> fn(&'a i32) -> &'a i32) = &((|x| x) as
fn(&_) -> &_);
let supertype: &(fn(&'static i32) -> &'static i32) = subtype;
// This works similarly for trait objects
let subtype: &(dyn for<'a> Fn(&'a i32) -> &'a i32) = &|x| x;
let supertype: &(dyn Fn(&'static i32) -> &'static i32) =
subtype;
// We can also substitute one higher-ranked lifetime for
another
let subtype: &(for<'a, 'b> fn(&'a i32, &'b i32)) = &((|x, y|)
```

{}) as fn(&\_, &\_));
let supertype: &for<'c> fn(&'c i32, &'c i32) = subtype;

r[subtyping.variance]

### Variance

r[subtyping.variance.intro] Variance is a property that generic types have with respect to their arguments. A generic type's *variance* in a parameter is how the subtyping of the parameter affects the subtyping of the type.

r[subtyping.variance.covariant]

F<T> is *covariant* over T if T being a subtype of U implies that F<T> is a subtype of F<U> (subtyping "passes through")

r[subtyping.variance.contravariant]

F<T> is *contravariant* over T if T being a subtype of U implies that
 F<U> is a subtype of F<T>

r[subtyping.variance.invariant]

F<T> is *invariant* over T otherwise (no subtyping relation can be derived)

r[subtyping.variance.builtin-types] Variance of types is automatically determined as follows

Туре	Variance in 'a	Variance in T
&'а Т	covariant	covariant
&'a mut T	covariant	invariant
*const T		covariant
*mut T		invariant
[T] and [T; n]		covariant
fn() -> T		covariant
fn(T) -> ()		contravariant
<pre>std::cell::UnsafeCell<t></t></pre>		invariant

Туре	Variance in 'a	Variance in T
<pre>std::marker::PhantomData<t></t></pre>		covariant
dyn Trait <t> + 'a</t>	covariant	invariant

r[subtyping.variance.user-composite-types] The variance of other struct, enum, and union types is decided by looking at the variance of the types of their fields. If the parameter is used in positions with different variances then the parameter is invariant. For example the following struct is covariant in 'a and T and invariant in 'b, 'c, and U.

```
use std::cell::UnsafeCell:
struct Variance<'a, 'b, 'c, T, U: 'a> {
                           // This makes `Variance` covariant
    x: &'a U,
in 'a, and would
                             // make it covariant in U, but U
is used later
   y: *const T,
                          // Covariant in T
    z: UnsafeCell<&'b f64>, // Invariant in 'b
    w: *mut U,
                          // Invariant in U, makes the whole
struct invariant
     f: fn(&'c ()) -> &'c () // Both co- and contravariant,
makes 'c invariant
                            // in the struct.
}
```

r[subtyping.variance.builtin-composite-types] When used outside of an struct, enum, or union, the variance for parameters is checked at each location separately.

```
# use std::cell::UnsafeCell;
fn generic tuple<'short, 'long: 'short>(
```

 $\ensuremath{//}$  'long is used inside of a tuple in both a co- and invariant position.

```
x: (&'long u32, UnsafeCell<&'long u32>),
) {
      // As the variance at these positions is computed
separately,
   // we can freely shrink 'long in the covariant position.
   let _: (&'short u32, UnsafeCell<&'long u32>) = x;
}
fn takes_fn_ptr<'short, 'middle: 'short>(
      // 'middle is used in both a co- and contravariant
position.
   f: fn(&'middle ()) -> &'middle (),
) {
      // As the variance at these positions is computed
separately,
   // we can freely shrink 'middle in the covariant position
   // and extend it in the contravariant position.
   let _: fn(&'static ()) -> &'short () = f;
}
```

r[bound]

## **Trait and lifetime bounds**

```
r[bound.syntax]
TypeParamBounds -> TypeParamBound ( `+` TypeParamBound )* `+`?
TypeParamBound -> Lifetime | TraitBound | UseBound
TraitBound ->
      (`?` | ForLifetimes )? TypePath
    | `(` ( `?` | ForLifetimes )? TypePath `)`
LifetimeBounds -> ( Lifetime `+` )* Lifetime?
Lifetime ->
      LIFETIME_OR_LABEL
    | `'static`
    | ` ' `
UseBound -> `use` UseBoundGenericArgs
UseBoundGenericArgs ->
      `<``>`
     | `<` ( UseBoundGenericArg `,`)* UseBoundGenericArg `,`?</pre>
`>`
UseBoundGenericArg ->
      Lifetime
    | IDENTIFIER
    | `Self`
```

r[bound.intro] <u>Trait</u> and lifetime bounds provide a way for <u>generic items</u> to restrict which types and lifetimes are used as their parameters. Bounds can be provided on any type in a <u>where clause</u>. There are also shorter forms for certain common cases:

- Bounds written after declaring a <u>generic parameter</u>: fn f<A: Copy>()
   {} is the same as fn f<A>() where A: Copy {}.
- In trait declarations as <u>supertraits</u>: trait Circle : Shape {} is equivalent to trait Circle where Self : Shape {}.
- In trait declarations as bounds on <u>associated types</u>: trait A { type
   B: Copy; } is equivalent to trait A where Self::B: Copy { type
   B; }.

r[bound.satisfaction] Bounds on an item must be satisfied when using the item. When type checking and borrow checking a generic item, the bounds can be used to determine that a trait is implemented for a type. For example, given Ty: Trait

- In the body of a generic function, methods from Trait can be called on Ty values. Likewise associated constants on the Trait can be used.
- Associated types from Trait can be used.
- Generic functions and types with a T: Trait bounds can be used with Ty being used for T.

```
# type Surface = i32;
trait Shape {
    fn draw(&self, surface: Surface);
    fn name() -> &'static str;
}
fn draw_twice<T: Shape>(surface: Surface, sh: T) {
    sh.draw(surface); // Can call method because T:
Shape
    sh.draw(surface);
}
fn copy_and_draw_twice<T: Copy>(surface: Surface, sh: T) where
T: Shape {
    let shape_copy = sh; // doesn't move sh because T:
Copy
```

```
draw_twice(surface, sh); // Can use generic function
because T: Shape
}
struct Figure<S: Shape>(S, S);
fn name_figure<U: Shape>(
    figure: Figure<U>, // Type Figure<U> is well-
formed because U: Shape
) {
    println!(
        "Figure of two {}",
        U::name(), // Can use associated function
    );
}
```

r[bound.trivial] Bounds that don't use the item's parameters or <u>higher-</u> <u>ranked lifetimes</u> are checked when the item is defined. It is an error for such a bound to be false.

r[bound.special] [Copy], [Clone], and [Sized] bounds are also checked for certain generic types when using the item, even if the use does not provide a concrete type. It is an error to have Copy or Clone as a bound on a mutable reference, <u>trait object</u>, or <u>slice</u>. It is an error to have Sized as a bound on a trait object or slice.

```
struct A<'a, T>
where
    i32: Default, // Allowed, but not useful
    i32: Iterator, // Error: `i32` is not an iterator
    &'a mut T: Copy, // (at use) Error: the trait bound
is not satisfied
    [T]: Sized, // (at use) Error: size cannot be
known at compilation
{
    f: &'a T,
```

```
}
struct UsesA<'a, T>(A<'a, T>);
```

r[bound.trait-object] Trait and lifetime bounds are also used to name <u>trait</u> <u>objects</u>.

r[bound.sized]

### ?Sized

? is only used to relax the implicit [Sized] trait bound for <u>type</u> <u>parameters</u> or <u>associated types</u>. ?Sized may not be used as a bound for other types.

r[bound.lifetime]

### **Lifetime bounds**

r[bound.lifetime.intro] Lifetime bounds can be applied to types or to other lifetimes.

r[bound.lifetime.outlive-lifetime] The bound 'a: 'b is usually read as 'a *outlives* 'b. 'a: 'b means that 'a lasts at least as long as 'b, so a reference &'a () is valid whenever &'b () is valid.

r[bound.lifetime.outlive-type] T: 'a means that all lifetime parameters
of T outlive 'a. For example, if 'a is an unconstrained lifetime parameter,
then i32: 'static and &'static str: 'a are satisfied, but Vec<&'a
()>: 'static is not.

r[bound.higher-ranked]

## **Higher-ranked trait bounds**

```
r[bound.higher-ranked.syntax]
ForLifetimes -> `for` GenericParams
```

```
r[bound.higher-ranked.intro] Trait bounds may be higher ranked over lifetimes. These bounds specify a bound that is true for all lifetimes. For example, a bound such as for<'a> &'a T: PartialEq<i32> would require an implementation like
```

and could then be used to compare a &'a T with any lifetime to an i32.

Only a higher-ranked bound can be used here, because the lifetime of the reference is shorter than any possible lifetime parameter on the function:

```
fn call_on_ref_zero<F>(f: F) where for<'a> F: Fn(&'a i32) {
    let zero = 0;
    f(&zero);
}
```

r[bound.higher-ranked.trait] Higher-ranked lifetimes may also be specified just before the trait: the only difference is the <u>scope</u> of the lifetime parameter, which extends only to the end of the following trait instead of the whole bound. This function is equivalent to the last one.

```
fn call_on_ref_zero<F>(f: F) where F: for<'a> Fn(&'a i32) {
    let zero = 0;
    f(&zero);
}
```

r[bound.implied]

## **Implied bounds**

r[bound.implied.intro] Lifetime bounds required for types to be wellformed are sometimes inferred.

```
fn requires_t_outlives_a<'a, T>(x: &'a T) {}
```

The type parameter  $\top$  is required to outlive 'a for the type &'a  $\top$  to be well-formed. This is inferred because the function signature contains the type &'a  $\top$  which is only valid if  $\top$ : 'a holds.

r[bound.implied.context] Implied bounds are added for all parameters and outputs of functions. Inside of requires\_t\_outlives\_a you can assume T: 'a to hold even if you don't explicitly specify this:

```
fn requires_t_outlives_a_not_implied<'a, T: 'a>() {}
fn requires_t_outlives_a<'a, T>(x: &'a T) {
    // This compiles, because `T: 'a` is implied by
    // the reference type `&'a T`.
    requires_t_outlives_a_not_implied::<'a, T>();
}
# fn requires_t_outlives_a_not_implied<'a, T: 'a>() {}
fn not_implied<'a, T>() {
    // This errors, because `T: 'a` is not implied by
    // the function signature.
    requires_t_outlives_a_not_implied::<'a, T>();
}
```

r[bound.implied.trait] Only lifetime bounds are implied, trait bounds still have to be explicitly added. The following example therefore causes an error:

```
use std::fmt::Debug;
struct IsDebug<T: Debug>(T);
// error[E0277]: `T` doesn't implement `Debug`
fn doesnt_specify_t_debug<T>(x: IsDebug<T>) {}
```

r[bound.implied.def] Lifetime bounds are also inferred for type definitions and impl blocks for any type:

```
struct Struct<'a, T> {
    // This requires `T: 'a` to be well-formed
    // which is inferred by the compiler.
    field: &'a T,
}
enum Enum<'a, T> {
    // This requires `T: 'a` to be well-formed,
    // which is inferred by the compiler.
    11
    // Note that `T: 'a` is required even when only
    // using `Enum::OtherVariant`.
    SomeVariant(&'a T),
    OtherVariant,
}
trait Trait<'a, T: 'a> {}
// This would error because `T: 'a` is not implied by any type
// in the impl header.
       impl<'a, T> Trait<'a, T> for () {}
//
// This compiles as `T: 'a` is implied by the self type `&'a
Τ`.
impl<'a, T> Trait<'a, T> for &'a T {}
 r[bound.use]
```

### **Use bounds**

Certain bounds lists may include a use<..> bound to control which generic parameters are captured by the impl Trait abstract return type. See precise capturing for more details.

r[coerce]

## **Type coercions**

r[coerce.intro] **Type coercions** are implicit operations that change the type of a value. They happen automatically at specific locations and are highly restricted in what types actually coerce.

r[coerce.as] Any conversions allowed by coercion can also be explicitly performed by the <u>type cast operator</u>, as.

Coercions are originally defined in <u>RFC 401</u> and expanded upon in <u>RFC 1558</u>.

r[coerce.site]

## **Coercion sites**

r[coerce.site.intro] A coercion can only occur at certain coercion sites in a program; these are typically places where the desired type is explicit or can be derived by propagation from explicit types (without type inference). Possible coercion sites are:

r[coerce.site.let]

• let statements where an explicit type is given.

For example, &mut 42 is coerced to have type &i8 in the following:

let \_: &i8 = &mut 42;

r[coerce.site.value]

• static and const item declarations (similar to let statements).

r[coerce.site.argument]

• Arguments for function calls

The value being coerced is the actual parameter, and it is coerced to the type of the formal parameter.

For example, &mut 42 is coerced to have type &i8 in the following:

```
fn bar(_: &i8) { }
fn main() {
    bar(&mut 42);
}
```

For method calls, the receiver (self parameter) type is coerced differently, see the documentation on <u>method-call expressions</u> for details.

r[coerce.site.constructor]

• Instantiations of struct, union, or enum variant fields

For example, &mut 42 is coerced to have type &i8 in the following:

```
struct Foo<'a> { x: &'a i8 }
fn main() {
    Foo { x: &mut 42 };
}
```

r[coerce.site.return]

• Function results—either the final line of a block if it is not semicolon-terminated or any expression in a return statement

For example,  $\times$  is coerced to have type &dyn Display in the following:

r[coerce.site.subexpr] If the expression in one of these coercion sites is a coercion-propagating expression, then the relevant sub-expressions in that expression are also coercion sites. Propagation recurses from these new coercion sites. Propagating expressions and their relevant sub-expressions are:

r[coerce.site.array]

• Array literals, where the array has type [U; n]. Each sub-expression in the array literal is a coercion site for coercion to type U.

r[coerce.site.repeat]

Array literals with repeating syntax, where the array has type [U; n].
 The repeated sub-expression is a coercion site for coercion to type U.

r[coerce.site.tuple]
Tuples, where a tuple is a coercion site to type (U\_0, U\_1, ..., U\_n). Each sub-expression is a coercion site to the respective type, e.g. the zeroth sub-expression is a coercion site to type U\_0.

r[coerce.site.parenthesis]

Parenthesized sub-expressions ((e)): if the expression has type U, then the sub-expression is a coercion site to U.

r[coerce.site.block]

Blocks: if a block has type U, then the last expression in the block (if it is not semicolon-terminated) is a coercion site to U. This includes blocks which are part of control flow statements, such as if/else, if the block has a known type.

r[coerce.types]

## **Coercion types**

r[coerce.types.intro] Coercion is allowed between the following types: r[coerce.types.reflexive]

- T to U if T is a <u>subtype</u> of U (*reflexive case*) r[coerce.types.transitive]
- r[coerce.types.transitive]
- T\_1 to T\_3 where T\_1 coerces to T\_2 and T\_2 coerces to T\_3 (*transitive case*)

Note that this is not fully supported yet. r[coerce.types.mut-reborrow]

• &mut T to &T

r[coerce.types.mut-pointer]

• \*mut T to \*const T r[coerce.types.ref-to-pointer]

• &T to \*const T

r[coerce.types.mut-to-pointer]

• &mut T to \*mut T

r[coerce.types.deref]

• &T or &mut T to &U if T implements Deref<Target = U>. For example:

```
use std::ops::Deref;
struct CharContainer {
    value: char,
}
impl Deref for CharContainer {
```

```
type Target = char;
fn deref<'a>(&'a self) -> &'a char {
    &self.value
  }
}
fn foo(arg: &char) {}
fn main() {
    let x = &mut CharContainer { value: 'y' };
    foo(x); //&mut CharContainer is coerced to &char.
}
```

r[coerce.types.deref-mut]

&mut T to &mut U if T implements DerefMut<Target = U>.
 r[coerce.types.unsize]

• TyCtor( $\top$ ) to TyCtor( $\cup$ ), where TyCtor( $\top$ ) is one of

```
    &T
    &mut T
    *const T
    *mut T
```

• Box<T>

and where  $\cup$  can be obtained from  $\top$  by <u>unsized coercion</u>. r[coerce.types.fn]

```
• Function item types to fn pointers r[coerce.types.closure]
```

```
• Non capturing closures to fn pointers r[coerce.types.never]
```

```
• ! to any T
r[coerce.unsize]
```

## **Unsized Coercions**

r[coerce.unsize.intro] The following coercions are called unsized coercions, since they relate to converting types to unsized types, and are permitted in a few cases where other coercions are not, as described above. They can still happen anywhere else a coercion can occur.

r[coerce.unsize.trait] Two traits, <u>Unsize</u> and <u>CoerceUnsized</u>, are used to assist in this process and expose it for library use. The following coercions are built-ins and, if T can be coerced to U with one of them, then an implementation of <u>Unsize<U></u> for T will be provided:

r[coerce.unsize.slice]

```
• [T; n] to [T].
```

r[coerce.unsize.trait-object]

• T to dyn U, when T implements U + Sized, and U is <u>dyn</u> <u>compatible</u>.

r[coerce.unsize.trait-upcast]

- dyn T to dyn U, when U is one of T's <u>supertraits</u>.
  - This allows dropping auto traits, i.e. dyn T + Auto to dyn U is allowed.
  - This allows adding auto traits if the principal trait has the auto trait as a super trait, i.e. given trait T: U + Send {}, dyn T to dyn T + Send or to dyn U + Send coercions are allowed.

r[coerce.unsized.composite]

- Foo<..., T, ...> to Foo<..., U, ...>, when:
  - Foo is a struct.
  - T implements Unsize<U>.

- The last field of Foo has a type involving T.
- If that field has type Bar<T>, then Bar<T> implements Unsize<Bar<U>>.
- T is not part of the type of any other fields.

r[coerce.unsized.pointer] Additionally, a type Foo<T> can implement CoerceUnsized<Foo<U>> when T implements Unsize<U> or CoerceUnsized<Foo<U>>. This allows it to provide an unsized coercion to Foo<U>.

[!NOTE] While the definition of the unsized coercions and their implementation has been stabilized, the traits themselves are not yet stable and therefore can't be used directly in stable Rust.

r[coerce.least-upper-bound]

# **Least upper bound coercions**

r[coerce.least-upper-bound.intro] In some contexts, the compiler must coerce together multiple types to try and find the most general type. This is called a "Least Upper Bound" coercion. LUB coercion is used and only used in the following situations:

- To find the common type for a series of if branches.
- To find the common type for a series of match arms.
- To find the common type for array elements.
- To find the type for the return type of a closure with multiple return statements.
- To check the type for the return type of a function with multiple return statements.

r[coerce.least-upper-bound.target] In each such case, there are a set of types T0..Tn to be mutually coerced to some target type  $T_t$ , which is unknown to start.

r[coerce.least-upper-bound.computation] Computing the LUB coercion is done iteratively. The target type  $T_t$  begins as the type  $T_0$ . For each new type  $T_i$ , we consider whether

r[coerce.least-upper-bound.computation-identity]

• If Ti can be coerced to the current target type T\_t, then no change is made.

r[coerce.least-upper-bound.computation-replace]

- Otherwise, check whether T\_t can be coerced to Ti; if so, the T\_t is changed to Ti. (This check is also conditioned on whether all of the source expressions considered thus far have implicit coercions.)
   r[coerce.least-upper-bound.computation-unify]
- If not, try to compute a mutual supertype of T\_t and Ti, which will become the new target type.

# **Examples:**

```
# let (a, b, c) = (0, 1, 2);
// For if branches
let bar = if true {
   а
} else if false {
  b
} else {
  С
};
// For match arms
let baw = match 42 {
   0 => a,
   1 => b,
   _ => c,
};
// For array elements
let bax = [a, b, c];
// For closure with multiple return statements
let clo = || {
   if true {
        а
    } else if false {
       b
   } else {
       С
    }
};
let baz = clo();
// For type checking of function with multiple return
statements
fn foo() -> i32 {
```

```
let (a, b, c) = (0, 1, 2);
match 42 {
        0 => a,
        1 => b,
        _ => c,
   }
}
```

In these examples, types of the ba<sup>\*</sup> are found by LUB coercion. And the compiler checks whether LUB coercion result of a, b, c is i32 in the processing of the function foo.

### Caveat

This description is obviously informal. Making it more precise is expected to proceed as part of a general effort to specify the Rust type checker more precisely. r<u>destructors</u>

# **Destructors**

r[destructors.intro] When an <u>initialized variable</u> or <u>temporary</u> goes out of <u>scope</u>, its *destructor* is run, or it is *dropped*. <u>Assignment</u> also runs the destructor of its left-hand operand, if it's initialized. If a variable has been partially initialized, only its initialized fields are dropped.

r[destructors.operation] The destructor of a type T consists of:

1. If T: Drop, calling <<u>T</u> as std::ops::Drop>::drop

```
2. Recursively running the destructor of all of its fields.
```

- The fields of a <u>struct</u> are dropped in declaration order.
- The fields of the active <u>enum variant</u> are dropped in declaration order.
- The fields of a <u>tuple</u> are dropped in order.
- The elements of an <u>array</u> or owned <u>slice</u> are dropped from the first element to the last.
- The variables that a <u>closure</u> captures by move are dropped in an unspecified order.
- <u>Trait objects</u> run the destructor of the underlying type.
- Other types don't result in any further drops.

r[destructors.drop\_in\_place] If a destructor must be run manually, such as when implementing your own smart pointer, [std::ptr::drop\_in\_place] can be used.

Some examples:

```
struct PrintOnDrop(&'static str);
impl Drop for PrintOnDrop {
    fn drop(&mut self) {
        println!("{}", self.0);
     }
}
let mut overwritten = PrintOnDrop("drops when overwritten");
```

```
overwritten = PrintOnDrop("drops when scope ends");
let tuple = (PrintOnDrop("Tuple first"), PrintOnDrop("Tuple
second"));
let moved;
// No destructor run on assignment.
moved = PrintOnDrop("Drops when moved");
// Drops now, but is then uninitialized.
moved;
// Uninitialized does not drop.
let uninitialized: PrintOnDrop;
// After a partial move, only the remaining fields are
dropped.
               partial_move = (PrintOnDrop("first"),
let
       mut
PrintOnDrop("forgotten"));
   Perform a partial move, leaving only `partial_move.0`
11
initialized.
core::mem::forget(partial_move.1);
// When partial_move's scope ends, only the first field is
dropped.
```

```
r[destructors.scope]
```

## **Drop scopes**

r[destructors.scope.intro] Each variable or temporary is associated to a *drop scope*. When control flow leaves a drop scope all variables associated to that scope are dropped in reverse order of declaration (for variables) or creation (for temporaries).

r[destructors.scope.desugaring] Drop scopes can be determined by replacing <u>for</u>, <u>if</u>, and <u>while</u> expressions with equivalent expressions using <u>match</u>, <u>loop</u> and break.

r[destructors.scope.operators] Overloaded operators are not distinguished from built-in operators and <u>binding modes</u> are not considered.

r[destructors.scope.list] Given a function, or closure, there are drop scopes for:

r[destructors.scope.function]

- The entire function r[destructors.scope.statement]
- Each <u>statement</u> r[destructors.scope.expression]
- Each <u>expression</u> r[destructors.scope.block]
- Each block, including the function body
  - In the case of a <u>block expression</u>, the scope for the block and the expression are the same scope.

r[destructors.scope.match-arm]

• Each arm of a match expression

r[destructors.scope.nesting] Drop scopes are nested within one another as follows. When multiple scopes are left at once, such as when returning from a function, variables are dropped from the inside outwards. r[destructors.scope.nesting.function]

• The entire function scope is the outer most scope. r[destructors.scope.nesting.function-body]

• The function body block is contained within the scope of the entire function.

r[destructors.scope.nesting.expr-statement]

• The parent of the expression in an expression statement is the scope of the statement.

r[destructors.scope.nesting.let-initializer]

• The parent of the initializer of a <u>let statement</u> is the <u>let</u> statement's scope.

r[destructors.scope.nesting.statement]

• The parent of a statement scope is the scope of the block that contains the statement.

r[destructors.scope.nesting.match-guard]

• The parent of the expression for a match guard is the scope of the arm that the guard is for.

r[destructors.scope.nesting.match-arm]

• The parent of the expression after the => in a match expression is the scope of the arm that it's in.

r[destructors.scope.nesting.match]

• The parent of the arm scope is the scope of the match expression that it belongs to.

r[destructors.scope.nesting.other]

• The parent of all other scopes is the scope of the immediately enclosing expression.

r[destructors.scope.params]

## **Scopes of function parameters**

All function parameters are in the scope of the entire function body, so are dropped last when evaluating the function. Each actual function parameter is dropped after any bindings introduced in that parameter's pattern.

```
# struct PrintOnDrop(&'static str);
# impl Drop for PrintOnDrop {
      fn drop(&mut self) {
#
          println!("drop({})", self.0);
#
#
      }
# }
// Drops `y`, then the second parameter, then `x`, then the
first parameter
fn patterns_in_parameters(
    (x, _): (PrintOnDrop, PrintOnDrop),
    (_, y): (PrintOnDrop, PrintOnDrop),
) {}
// drop order is 3 2 0 1
patterns_in_parameters(
    (PrintOnDrop("0"), PrintOnDrop("1")),
    (PrintOnDrop("2"), PrintOnDrop("3")),
);
```

r[destructors.scope.bindings]

# Scopes of local variables

r[destructors.scope.bindings.intro] Local variables declared in a let statement are associated to the scope of the block that contains the let statement. Local variables declared in a match expression are associated to the arm scope of the match arm that they are declared in.

```
# struct PrintOnDrop(&'static str);
# impl Drop for PrintOnDrop {
```

```
fn drop(&mut self) {
#
          println!("drop({})", self.0);
#
#
      }
# }
     declared_first = PrintOnDrop("Dropped
let
                                               last
                                                     in
                                                          outer
scope");
{
      let declared_in_block = PrintOnDrop("Dropped in inner
scope");
}
let
     declared_last = PrintOnDrop("Dropped
                                              first
                                                     in
                                                          outer
scope");
```

r[destructors.scope.bindings.match-pattern-order] If multiple patterns are used in the same arm for a match expression, then an unspecified pattern will be used to determine the drop order.

r[destructors.scope.temporary]

## **Temporary scopes**

r[destructors.scope.temporary.intro] The *temporary scope* of an expression is the scope that is used for the temporary variable that holds the result of that expression when used in a <u>place context</u>, unless it is <u>promoted</u>.

r[destructors.scope.temporary.enclosing] Apart from lifetime extension, the temporary scope of an expression is the smallest scope that contains the expression and is one of the following:

- The entire function.
- A statement.
- The body of an <u>if</u>, <u>while</u> or <u>loop</u> expression.
- The else block of an if expression.
- The non-pattern matching condition expression of an if or while expression, or a match guard.
- The body expression for a match arm.
- Each operand of a <u>lazy boolean expression</u>.

- The pattern-matching condition(s) and consequent body of <u>if</u> ([destructors.scope.temporary.edition2024]).
- The pattern-matching condition and loop body of while.
- The entirety of the tail expression of a block ([destructors.scope.temporary.edition2024]).

[!NOTE] The <u>scrutinee</u> of a match expression is not a temporary scope, so temporaries in the scrutinee can be dropped after the match expression. For example, the temporary for 1 in match 1 { ref mut  $z \Rightarrow z$  }; lives until the end of the statement.

#### r[destructors.scope.temporary.edition2024]

[!EDITION-2024] The 2024 edition added two new temporary scope narrowing rules: if let temporaries are dropped before the else block, and temporaries of tail expressions of blocks are dropped immediately after the tail expression is evaluated.

Some examples:

```
# #![allow(irrefutable_let_patterns)]
# struct PrintOnDrop(&'static str);
# impl Drop for PrintOnDrop {
      fn drop(&mut self) {
#
          println!("drop({})", self.0);
#
      }
#
# }
let local_var = PrintOnDrop("local var");
// Dropped once the condition has been evaluated
if PrintOnDrop("If condition").0 == "If condition" {
    // Dropped at the end of the block
    PrintOnDrop("If body").0
} else {
    unreachable!()
};
if let "if let scrutinee" = PrintOnDrop("if let scrutinee").0
```

```
{
    PrintOnDrop("if let consequent").0
    // `if let consequent` dropped here
}
// `if let scrutinee` is dropped here
else {
    PrintOnDrop("if let else").0
    // `if let else` dropped here
};
while let x = PrintOnDrop("while let scrutinee").0 {
    PrintOnDrop("while let loop body").0;
    break;
    // `while let loop body` dropped here.
    // `while let scrutinee` dropped here.
}
// Dropped before the first ||
(PrintOnDrop("first operand").0 == ""
// Dropped before the )
|| PrintOnDrop("second operand").0 == "")
// Dropped before the ;
|| PrintOnDrop("third operand").0 == "";
// Scrutinee is dropped at the end of the function, before
local variables
// (because this is the tail expression of the function body
block).
match PrintOnDrop("Matched value in final expression") {
    // Dropped once the condition has been evaluated
   _ if PrintOnDrop("guard condition").0 == "" => (),
   _ => (),
}
```

```
r[destructors.scope.operands]
```

# Operands

Temporaries are also created to hold the result of operands to an expression while the other operands are evaluated. The temporaries are associated to the scope of the expression with that operand. Since the temporaries are moved from once the expression is evaluated, dropping them has no effect unless one of the operands to an expression breaks out of the expression, returns, or <u>panics</u>.

```
# struct PrintOnDrop(&'static str);
# impl Drop for PrintOnDrop {
#
      fn drop(&mut self) {
          println!("drop({})", self.0);
#
#
      }
# }
loop {
    // Tuple expression doesn't finish evaluating so operands
drop in reverse order
    (
        PrintOnDrop("Outer tuple first"),
        PrintOnDrop("Outer tuple second"),
        (
            PrintOnDrop("Inner tuple first"),
            PrintOnDrop("Inner tuple second"),
            break,
        ),
        PrintOnDrop("Never created"),
    );
}
```

r[destructors.scope.const-promotion]

### **Constant promotion**

Promotion of a value expression to a 'static slot occurs when the expression could be written in a constant and borrowed, and that borrow could be dereferenced where the expression was originally written, without changing the runtime behavior. That is, the promoted expression can be

evaluated at compile-time and the resulting value does not contain <u>interior</u> <u>mutability</u> or <u>destructors</u> (these properties are determined based on the value where possible, e.g. <u>&None</u> always has the type <u>&'static</u> Option<\_>, as it contains nothing disallowed).

r[destructors.scope.lifetime-extension]

### **Temporary lifetime extension**

[!NOTE] The exact rules for temporary lifetime extension are subject to change. This is describing the current behavior only.

r[destructors.scope.lifetime-extension.let] The temporary scopes for expressions in let statements are sometimes *extended* to the scope of the block containing the let statement. This is done when the usual temporary scope would be too small, based on certain syntactic rules. For example:

```
let x = &mut 0;
// Usually a temporary would be dropped by now, but the
temporary for `0` lives
// to the end of the block.
println!("{}", x);
```

r[destructors.scope.lifetime-extension.static] Lifetime extension also applies to static and const items, where it makes temporaries live until the end of the program. For example:

```
const C: &Vec<i32> = &Vec::new();
// Usually this would be a dangling reference as the `Vec`
would only
// exist inside the initializer expression of `C`, but instead
the
// borrow gets lifetime-extended so it effectively has
`'static` lifetime.
println!("{:?}", C);
```

r[destructors.scope.lifetime-extension.sub-expressions] If a <u>borrow</u>, <u>dereference</u>, <u>field</u>, or <u>tuple indexing expression</u> has an extended temporary scope then so does its operand. If an <u>indexing expression</u> has an extended

temporary scope then the indexed expression also has an extended temporary scope.

r[destructors.scope.lifetime-extension.patterns]

#### **Extending based on patterns**

r[destructors.scope.lifetime-extension.patterns.extending] An *extending pattern* is either

- An <u>identifier pattern</u> that binds by reference or mutable reference.
- A <u>struct</u>, <u>tuple</u>, <u>tuple struct</u>, or <u>slice</u> pattern where at least one of the direct subpatterns is an extending pattern.

```
So ref x, V(ref x) and [ref x, y] are all extending patterns, but x, &ref x and &(ref x,) are not.
```

r[destructors.scope.lifetime-extension.patterns.let] If the pattern in a let statement is an extending pattern then the temporary scope of the initializer expression is extended.

r[destructors.scope.lifetime-extension.exprs]

#### **Extending based on expressions**

For a let statement with an initializer, an *extending expression* is an expression which is one of the following:

- The initializer expression.
- The operand of an extending <u>borrow expression</u>.
- The operand(s) of an extending <u>array</u>, <u>cast</u>, <u>braced struct</u>, or <u>tuple</u> expression.
- The arguments to an extending <u>tuple struct</u> or <u>tuple variant</u> constructor expression.
- The final expression of any extending <u>block expression</u>.

So the borrow expressions in &mut 0, (&1, &mut 2), and Some(&mut 3) are all extending expressions. The borrows in &0 + &1 and f(&mut 0) are not.

The operand of any extending borrow expression has its temporary scope extended.

#### Examples

Here are some examples where expressions have extended temporary scopes:

```
# fn temp() {}
# trait Use { fn use_temp(&self) -> &Self { self } }
# impl Use for () {}
// The temporary that stores the result of `temp()` lives in
the same scope
// as x in these cases.
let x = &temp();
let x = &temp() as &dyn Send;
let x = (&*&temp(),);
let x = { [Some(&temp()) ] };
let ref x = temp();
let ref x = *&temp();
# x;
```

Here are some examples where expressions don't have extended temporary scopes:

```
# fn temp() {}
# trait Use { fn use_temp(&self) -> &Self { self } }
# impl Use for () {}
// The temporary that stores the result of `temp()` only lives
until the
// end of the let statement in these cases.
let x = std::convert::identity(&temp()); // ERROR
let x = (&temp()).use_temp(); // ERROR
# x;
r[destructors.forget]
```

# Not running destructors

r[destructors.manually-suppressing]

## Manually suppressing destructors

[std::mem::forget] can be used to prevent the destructor of a variable from being run, and [std::mem::ManuallyDrop] provides a wrapper to prevent a variable or field from being dropped automatically.

[!NOTE] Preventing a destructor from being run via [std::mem::forget] or other means is safe even if it has a type that isn't 'static. Besides the places where destructors are guaranteed to run as defined by this document, types may *not* safely rely on a destructor being run for soundness.

r[destructors.process-termination]

## **Process termination without unwinding**

There are some ways to terminate the process without <u>unwinding</u>, in which case destructors will not be run.

The standard library provides [std::process::exit] and [std::process::abort] to do this explicitly. Additionally, if the [panic handler][panic.panic\_handler.std] is set to abort, panicking will always terminate the process without destructors being run.

There is one additional case to be aware of: when a panic reaches a <u>non-unwinding ABI boundary</u>, either no destructors will run, or all destructors up until the ABI boundary will run.

r[lifetime-elision]

# Lifetime elision

Rust has rules that allow lifetimes to be elided in various places where the compiler can infer a sensible default choice.

r[lifetime-elision.function]

# Lifetime elision in functions

r[lifetime-elision.function.intro] In order to make common patterns more ergonomic, lifetime arguments can be *elided* in <u>function item</u>, <u>function</u> <u>pointer</u>, and <u>closure trait</u> signatures. The following rules are used to infer lifetime parameters for elided lifetimes.

r[lifetime-elision.function.lifetimes-not-inferred] It is an error to elide lifetime parameters that cannot be inferred.

r[lifetime-elision.function.explicit-placeholder] The placeholder lifetime, '\_\_, can also be used to have a lifetime inferred in the same way. For lifetimes in paths, using '\_\_ is preferred.

r[lifetime-elision.function.only-functions] Trait object lifetimes follow different rules discussed <u>below</u>.

r[lifetime-elision.function.implicit-lifetime-parameters]

• Each elided lifetime in the parameters becomes a distinct lifetime parameter.

r[lifetime-elision.function.output-lifetime]

• If there is exactly one lifetime used in the parameters (elided or not), that lifetime is assigned to *all* elided output lifetimes.

r[lifetime-elision.function.receiver-lifetime] In method signatures there is another rule

• If the receiver has type &Self or &mut Self, then the lifetime of that reference to Self is assigned to all elided output lifetime parameters.

Examples:

```
# trait T {}
# trait ToCStr {}
# struct Thing<'a> {f: &'a i32}
# struct Command;
#
# trait Example {
fn print1(s: &str);
```

```
elided
fn print2(s: &'_ str);
                                                       // also
elided
fn print3<'a>(s: &'a str);
                                                            //
expanded
fn debug1(lvl: usize, s: &str);
                                                            11
elided
fn debug2<'a>(lvl: usize, s: &'a str);
                                                            11
expanded
fn substr1(s: &str, until: usize) -> &str;
                                                            11
elided
fn substr2<'a>(s: &'a str, until: usize) -> &'a str;
                                                           11
expanded
fn get_mut1(&mut self) -> &mut dyn T;
                                                            11
elided
fn get_mut2<'a>(&'a mut self) -> &'a mut dyn T;
                                                            11
expanded
fn args1<T: ToCStr>(&mut self, args: &[T]) -> &mut Command;
// elided
fn args2<'a, 'b, T: ToCStr>(&'a mut self, args: &'b [T]) ->
&'a mut Command; // expanded
fn other_args1<'a>(arg: &str) -> &'a str;
                                                            11
elided
fn other_args2<'a, 'b>(arg: &'b str) -> &'a str;
                                                            11
expanded
fn new1(buf: &mut [u8]) -> Thing<'_>;
                                                            //
elided - preferred
fn new2(buf: &mut [u8]) -> Thing;
                                                            11
elided
fn new3<'a>(buf: &'a mut [u8]) -> Thing<'a>;
                                                            //
```

```
expanded
 # }
 type FunPtr1 = fn(\&str) \rightarrow \&str;
                                                              11
 elided
 type FunPtr2 = for<'a> fn(&'a str) -> &'a str;
                                                              //
 expanded
 type FunTrait1 = dyn Fn(&str) -> &str;
                                                              //
 elided
 type FunTrait2 = dyn for<'a> Fn(&'a str) -> &'a str;
                                                              11
 expanded
// The following examples show situations where it is not
allowed to elide the
// lifetime parameter.
# trait Example {
// Cannot infer, because there are no parameters to infer from.
                                                               11
fn get_str() -> &str;
ILLEGAL
// Cannot infer, ambiguous if it is borrowed from the first or
second parameter.
fn frob(s: &str, t: &str) -> &str;
                                                               11
ILLEGAL
# }
  r[lifetime-elision.trait-object]
```

# **Default trait object lifetimes**

r[lifetime-elision.trait-object.intro] The assumed lifetime of references held by a <u>trait object</u> is called its *default object lifetime bound*. These were defined in <u>RFC 599</u> and amended in <u>RFC 1156</u>.

r[lifetime-elision.trait-object.explicit-bound] These default object lifetime bounds are used instead of the lifetime parameter elision rules defined above when the lifetime bound is omitted entirely.

r[lifetime-elision.trait-object.explicit-placeholder] If '\_\_\_\_ is used as the lifetime bound then the bound follows the usual elision rules.

r[lifetime-elision.trait-object.containing-type] If the trait object is used as a type argument of a generic type then the containing type is first used to try to infer a bound.

r[lifetime-elision.trait-object.containing-type-unique]

• If there is a unique bound from the containing type then that is the default

r[lifetime-elision.trait-object.containing-type-explicit]

• If there is more than one bound from the containing type then an explicit bound must be specified

r[lifetime-elision.trait-object.trait-bounds] If neither of those rules apply, then the bounds on the trait are used:

r[lifetime-elision.trait-object.trait-unique]

• If the trait is defined with a single lifetime *bound* then that bound is used.

r[lifetime-elision.trait-object.static-lifetime]

• If 'static is used for any lifetime bound then 'static is used. r[lifetime-elision.trait-object.default]

• If the trait has no lifetime bounds, then the lifetime is inferred in expressions and is 'static outside of expressions.

```
// For the following trait...
 trait Foo { }
 // These two are the same because Box<T> has no lifetime bound
 on T
 type T1 = Box<dyn Foo>;
 type T2 = Box<dyn Foo + 'static>;
 // ...and so are these:
 impl dyn Foo {}
 impl dyn Foo + 'static {}
 // ...so are these, because &'a T requires T: 'a
 type T3 < a > a > a dyn Foo;
 type T4 < a > = \& a (dyn Foo + a);
 // std::cell::Ref<'a, T> also requires T: 'a, so these are the
 same
 type T5<'a> = std::cell::Ref<'a, dyn Foo>;
 type T6<'a> = std::cell::Ref<'a, dyn Foo + 'a>;
// This is an example of an error.
# trait Foo { }
struct TwoBounds<'a, 'b, T: ?Sized + 'a + 'b> {
    f1: &'a i32,
    f2: &'b i32,
    f3: T,
}
type T7<'a, 'b> = TwoBounds<'a, 'b, dyn Foo>;
                                     ^^^^^
11
// Error: the lifetime bound for this object type cannot be
deduced from context
```

r[lifetime-elision.trait-object.innermost-type] Note that the innermost object sets the bound, so &'a Box<dyn Foo> is still &'a Box<dyn Foo + 'static>.

```
// For the following trait...
trait Bar<'a>: 'a { }
// ...these two are the same:
type T1<'a> = Box<dyn Bar<'a>>;
type T2<'a> = Box<dyn Bar<'a> + 'a>;
// ...and so are these:
impl<'a> dyn Bar<'a> {}
impl<'a> dyn Bar<'a> + 'a {}
```

```
r[lifetime-elision.const-static]
```

## const and static elision

r[lifetime-elision.const-static.implicit-static] Both <u>constant</u> and <u>static</u> declarations of reference types have *implicit* 'static lifetimes unless an explicit lifetime is specified. As such, the constant declarations involving 'static above may be written without the lifetimes.

```
// STRING: &'static str
const STRING: &str = "bitstring";
struct BitsNStrings<'a> {
    mybits: [u32; 2],
    mystring: &'a str,
}
// BITS_N_STRINGS: BitsNStrings<'static>
const BITS_N_STRINGS: BitsNStrings<'_> = BitsNStrings {
    mybits: [1, 2],
    mystring: STRING,
};
```

r[lifetime-elision.const-static.fn-references] Note that if the static or const items include function or closure references, which themselves include references, the compiler will first try the standard elision rules. If it is unable to resolve the lifetimes by its usual rules, then it will error. By way of example:

```
# struct Foo;
# struct Bar;
# struct Baz;
# fn somefunc(a: &Foo, b: &Bar, c: &Baz) -> usize {42}
// Resolved as `for<'a> fn(&'a str) -> &'a str`.
const RESOLVED_SINGLE: fn(&str) -> &str = |x| x;
// Resolved as `for<'a, 'b, 'c> Fn(&'a Foo, &'b Bar, &'c Baz)
-> usize`.
```

const RESOLVED\_MULTIPLE: &dyn Fn(&Foo, &Bar, &Baz) -> usize =
&somefunc;

# struct Foo; # struct Bar; # struct Baz; # fn somefunc<'a,'b>(a: &'a Foo, b: &'b Bar) -> &'a Baz {unimplemented!()}

// There is insufficient information to bound the return
reference lifetime

// relative to the argument lifetimes, so this is an error.

const RESOLVED\_STATIC: &dyn Fn(&Foo, &Bar) -> &Baz = &somefunc;
// ^

// this function's return type contains a borrowed value, but
the signature

// does not say whether it is borrowed from argument 1 or argument 2 r[lang-types]

# **Special types and traits**

r[lang-types.intro] Certain types and traits that exist in <u>the standard</u> <u>library</u> are known to the Rust compiler. This chapter documents the special features of these types and traits.

r[lang-types.box]

#### Box<T>

r[lang-types.box.intro] [Box<T>] has a few special features that Rust doesn't currently allow for user defined types.

r[lang-types.box.deref]

The <u>dereference operator</u> for Box<T> produces a place which can be moved from. This means that the \* operator and the destructor of Box<T> are built-in to the language.

r[lang-types.box.receiver]

- <u>Methods</u> can take Box<Self> as a receiver. r[lang-types.box.fundamental]
- A trait may be implemented for Box<T> in the same crate as T, which the <u>orphan rules</u> prevent for other generic types.
   r[lang-types.rc]

#### Rc<T>

r[lang-types.rc.receiver] Methods can take Rc<Self> as a receiver.
r[lang-types.arc]
## Arc<T>

r[lang-types.arc.receiver] Methods can take Arc<Self> as a receiver.
r[lang-types.pin]

#### Pin<P>

r[lang-types.pin.receiver] Methods can take Pin<P> as a receiver.
r[lang-types.unsafe-cell]

### UnsafeCell<T>

r[lang-types.unsafe-cell.interior-mut] [std::cell::UnsafeCell<T>] is used for <u>interior mutability</u>. It ensures that the compiler doesn't perform optimisations that are incorrect for such types.

r[lang-types.unsafe-cell.read-only-alloc] It also ensures that <u>static</u> <u>items</u> which have a type with interior mutability aren't placed in memory marked as read only.

r[lang-types.phantom-data]

#### PhantomData<T>

[std::marker::PhantomData<T>] is a zero-sized, minimum alignment, type that is considered to own a  $\top$  for the purposes of <u>variance</u>, <u>drop check</u>, and <u>auto traits</u>.

r[lang-types.ops]

# **Operator Traits**

The traits in [std::ops] and [std::cmp] are used to overload <u>operators</u>, <u>indexing expressions</u>, and <u>call expressions</u>.

r[lang-types.deref]

## Deref and DerefMut

As well as overloading the unary **\*** operator, <u>Deref</u> and <u>DerefMut</u> are also used in <u>method resolution</u> and <u>deref coercions</u>.

r[lang-types.drop]

### Drop

The [Drop] trait provides a <u>destructor</u>, to be run whenever a value of this type is to be destroyed.

r[lang-types.copy]

r[lang-types.copy.intro] The [Copy] trait changes the semantics of a type implementing it.

r[lang-types.copy.behavior] Values whose type implements Copy are copied rather than moved upon assignment.

r[lang-types.copy.constraint] Copy can only be implemented for types which do not implement Drop, and whose fields are all Copy. For enums, this means all fields of all variants have to be Copy. For unions, this means all variants have to be Copy.

r[lang-types.copy.builtin-types] Copy is implemented by the compiler for

r[lang-types.copy.tuple]

• <u>Tuples</u> of Copy types r[lang-types.copy.fn-pointer]

• <u>Function pointers</u> r[lang-types.copy.fn-item]

• <u>Function items</u> r[lang-types.copy.closure]

• <u>Closures</u> that capture no values or that only capture values of Copy types

r[lang-types.clone]

## Clone

r[lang-types.clone.intro] The [Clone] trait is a supertrait of Copy, so it also needs compiler generated implementations.

r[lang-types.clone.builtin-types] It is implemented by the compiler for the following types:

r[lang-types.clone.builtin-copy]

- Types with a built-in Copy implementation (see above) r[lang-types.clone.tuple]
- <u>Tuples</u> of Clone types

r[lang-types.clone.closure]

• <u>Closures</u> that only capture values of <u>Clone</u> types or capture no values from the environment

r[lang-types.send]

## Send

The [Send] trait indicates that a value of this type is safe to send from one thread to another.

r[lang-types.sync]

#### Sync

r[lang-types.sync.intro] The [Sync] trait indicates that a value of this type is safe to share between multiple threads.

r[lang-types.sync.static-constraint] This trait must be implemented for all types used in immutable <u>static items</u>.

r[lang-types.termination]

## Termination

The <u>Termination</u> trait indicates the acceptable return types for the <u>main</u> <u>function</u> and <u>test functions</u>.

r[lang-types.auto-traits]

## **Auto traits**

The [Send], [Sync], <u>Unpin</u>, <u>UnwindSafe</u>, and <u>RefUnwindSafe</u> traits are *auto traits*. Auto traits have special properties.

r[lang-types.auto-traits.auto-impl] If no explicit implementation or negative implementation is written out for an auto trait for a given type, then the compiler implements it automatically according to the following rules:

r[lang-types.auto-traits.builtin-composite]

• &T, &mut T, \*const T, \*mut T, [T; n], and [T] implement the trait if T does.

r[lang-types.auto-traits.fn-item-pointer]

• Function item types and function pointers automatically implement the trait.

r[lang-types.auto-traits.aggregate]

• Structs, enums, unions, and tuples implement the trait if all of their fields do.

r[lang-types.auto-traits.closure]

Closures implement the trait if the types of all of their captures do. A closure that captures a T by shared reference and a U by value implements any auto traits that both &T and U do.

r[lang-types.auto-traits.generic-impl] For generic types (counting the built-in types above as generic over T), if a generic implementation is available, then the compiler does not automatically implement it for types that could use the implementation except that they do not meet the requisite trait bounds. For instance, the standard library implements Send for all &T where T is Sync; this means that the compiler will not implement Send for &T if T is Send but not Sync.

 $\label{eq:relation} r[lang-types.auto-traits.negative] \mbox{ Auto traits can also have negative implementations, shown as implement for T in the standard for T is the standard for T in the standard for T in the standard for T is the standard for T in the standard for T in the standard for T is the standa$ 

library documentation, that override the automatic implementations. For example \*mut T has a negative implementation of Send, and so \*mut T is not Send, even if T is. There is currently no stable way to specify additional negative implementations; they exist only in the standard library.

r[lang-types.auto-traits.trait-object-marker] Auto traits may be added as an additional bound to any <u>trait object</u>, even though normally only one trait is allowed. For instance, Box<dyn Debug + Send + UnwindSafe> is a valid type.

r[lang-types.sized]

#### Sized

r[lang-types.sized.intro] The [Sized] trait indicates that the size of this type is known at compile-time; that is, it's not a <u>dynamically sized type</u>.

r[lang-types.sized.implicit-sized] <u>Type parameters</u> (except Self in traits) are Sized by default, as are <u>associated types</u>.

r[lang-types.sized.implicit-impl] Sized is always implemented automatically by the compiler, not by <u>implementation items</u>.

r[lang-types.sized.relaxation] These implicit Sized bounds may be relaxed by using the special ?Sized bound.

r[names]

# Names

r[names.intro] An *entity* is a language construct that can be referred to in some way within the source program, usually via a <u>path</u>. Entities include <u>types</u>, <u>items</u>, <u>generic parameters</u>, <u>variable bindings</u>, <u>loop labels</u>, <u>lifetimes</u>, <u>fields</u>, <u>attributes</u>, and <u>lints</u>.

r[names.decl] A *declaration* is a syntactical construct that can introduce a *name* to refer to an entity. Entity names are valid within a <u>scope</u> --- a region of source text where that name may be referenced.

r[names.explicit-decl] Some entities are <u>explicitly declared</u> in the source code, and some are <u>implicitly declared</u> as part of the language or compiler extensions.

r[names.path] <u>*Paths*</u> are used to refer to an entity, possibly in another module or type.

r[names.lifetime] Lifetimes and loop labels use a <u>dedicated syntax</u> using a leading quote.

r[names.namespace] Names are segregated into different <u>namespaces</u>, allowing entities in different namespaces to share the same name without conflict.

r[names.resolution] *<u>Name resolution</u>* is the compile-time process of tying paths, identifiers, and labels to entity declarations.

r[names.visibility] Access to certain names may be restricted based on their <u>visibility</u>.

r[names.explicit]

# **Explicitly declared entities**

r[names.explicit.list] Entities that explicitly introduce a name in the source code are:

r[names.explicit.item-decl]

- <u>Items</u>:
  - Module declarations
  - External crate declarations
  - <u>Use declarations</u>
  - <u>Function declarations</u> and <u>function parameters</u>
  - <u>Type aliases</u>
  - <u>struct</u>, <u>union</u>, <u>enum</u>, enum variant declarations, and their named fields
  - Constant item declarations
  - Static item declarations
  - <u>Trait item declarations</u> and their <u>associated items</u>
  - External block items
  - macro rules declarations and matcher metavariables
  - Implementation associated items

r[names.explicit.expr]

- <u>Expressions</u>:
  - <u>Closure</u> parameters
  - while let pattern bindings
  - for pattern bindings
  - <u>if let</u> pattern bindings
  - <u>match</u> pattern bindings
  - <u>Loop labels</u>

r[names.explicit.generics]

• <u>Generic parameters</u>

r[names.explicit.higher-ranked-bounds]

- <u>Higher ranked trait bounds</u> r[names.explicit.binding]
- <u>let statement</u> pattern bindings r[names.explicit.macro\_use]
- The <u>macro use attribute</u> can introduce macro names from another crate

r[names.explicit.macro\_export]

• The <u>macro export</u> <u>attribute</u> can introduce an alias for the macro into the crate root

r[names.explicit.macro-invocation] Additionally, <u>macro invocations</u> and <u>attributes</u> can introduce names by expanding to one of the above items.

r[names.implicit]

# **Implicitly declared entities**

r[names.implicit.list] The following entities are implicitly defined by the language, or are introduced by compiler options and extensions:

r[names.implicit.primitive-types]

- Language prelude:
  - <u>Boolean type</u> --- bool
  - Textual types --- char and str
  - Integer types --- i8, i16, i32, i64, i128, u8, u16, u32, u64, u128
  - Machine-dependent integer types --- usize and isize
  - <u>floating-point types</u> --- f32 and f64

r[names.implicit.builtin-attributes]

- <u>Built-in attributes</u> r[names.implicit.prelude]
- <u>Standard library prelude</u> items, attributes, and macros r[names.implicit.stdlib]
- <u>Standard library</u> crates in the root module r[names.implicit.extern-prelude]
- <u>External crates</u> linked by the compiler r[names.implicit.tool-attributes]
- <u>Tool attributes</u> r[names.implicit.lints]
- <u>Lints</u> and <u>tool lint attributes</u> r[names.implicit.derive-helpers]

• <u>Derive helper attributes</u> are valid within an item without being explicitly imported

r[names.implicit.lifetime-static]

• The <u>'static</u> lifetime

r[names.implicit.root] Additionally, the crate root module does not have a name, but can be referred to with certain <u>path qualifiers</u> or aliases.

r[names.namespaces]

# Namespaces

r[names.namespaces.intro] A *namespace* is a logical grouping of declared <u>names</u>. Names are segregated into separate namespaces based on the kind of entity the name refers to. Namespaces allow the occurrence of a name in one namespace to not conflict with the same name in another namespace.

There are several different namespaces that each contain different kinds of entities. The usage of a name will look for the declaration of that name in different namespaces, based on the context, as described in the <u>name</u> <u>resolution</u> chapter.

r[names.namespaces.kinds] The following is a list of namespaces, with their corresponding entities:

- Type Namespace
  - <u>Module declarations</u>
  - External crate declarations
  - External crate prelude items
  - <u>Struct</u>, <u>union</u>, <u>enum</u>, enum variant declarations
  - <u>Trait item declarations</u>
  - <u>Type aliases</u>
  - Associated type declarations
  - Built-in types: <u>boolean</u>, <u>numeric</u>, and <u>textual</u>
  - <u>Generic type parameters</u>
  - <u>Self</u> type
  - <u>Tool attribute modules</u>
- Value Namespace
  - Function declarations
  - <u>Constant item declarations</u>
  - Static item declarations
  - <u>Struct constructors</u>
  - Enum variant constructors

- <u>Self constructors</u>
- <u>Generic const parameters</u>
- Associated const declarations
- Associated function declarations
- Local bindings --- <u>let</u>, <u>if let</u>, <u>while let</u>, <u>for</u>, <u>match</u> arms, <u>function parameters</u>, <u>closure parameters</u>
- Captured <u>closure</u> variables
- Macro Namespace
  - <u>macro rules</u> declarations
  - Built-in attributes
  - <u>Tool attributes</u>
  - Function-like procedural macros
  - Derive macros
  - Derive macro helpers
  - <u>Attribute macros</u>
- Lifetime Namespace
  - Generic lifetime parameters
- Label Namespace
  - Loop labels
  - <u>Block labels</u>

An example of how overlapping names in different namespaces can be used unambiguously:

```
// Foo introduces a type in the type namespace and a
constructor in the value
// namespace.
struct Foo(u32);
// The `Foo` macro is declared in the macro namespace.
macro_rules! Foo {
   () => {};
}
```

```
// `Foo` in the `f` parameter type refers to `Foo` in the type
namespace.
// `'Foo` introduces a new lifetime in the lifetime namespace.
fn example<'Foo>(f: Foo) {
     // `Foo` refers to the `Foo` constructor in the value
namespace.
    let ctor = Foo;
    // `Foo` refers to the `Foo` macro in the macro namespace.
    Foo!{}
    // `'Foo` introduces a label in the label namespace.
    'Foo: loop {
        // `'Foo` refers to the `'Foo` lifetime parameter, and
`Foo`
        // refers to the type namespace.
        let x: &'Foo Foo;
        // `'Foo` refers to the label.
        break 'Foo;
    }
}
```

r[names.namespaces.without]

## Named entities without a namespace

The following entities have explicit names, but the names are not a part of any specific namespace.

# Fields

r[names.namespaces.without.fields] Even though struct, enum, and union fields are named, the named fields do not live in an explicit namespace. They can only be accessed via a <u>field expression</u>, which only inspects the field names of the specific type being accessed.

## **Use declarations**

r[names.namespaces.without.use] A <u>use declaration</u> has named aliases that it imports into scope, but the <u>use</u> item itself does not belong to a specific namespace. Instead, it can introduce aliases into multiple namespaces, depending on the item kind being imported.

r[names.namespaces.sub-namespaces]

## **Sub-namespaces**

r[names.namespaces.sub-namespaces.intro] The macro namespace is split into two sub-namespaces: one for <u>bang-style macros</u> and one for <u>attributes</u>. When an attribute is resolved, any bang-style macros in scope will be ignored. And conversely resolving a bang-style macro will ignore attribute macros in scope. This prevents one style from shadowing another.

For example, the <u>cfg\_attribute</u> and the <u>cfg\_macro</u> are two different entities with the same name in the macro namespace, but they can still be used in their respective context.

r[names.namespaces.sub-namespaces.use-shadow] It is still an error for a <u>use import</u> to shadow another macro, regardless of their sub-namespaces.

r[names.scopes]

# **Scopes**

r[names.scopes.intro] A *scope* is the region of source text where a named <u>entity</u> may be referenced with that name. The following sections provide details on the scoping rules and behavior, which depend on the kind of entity and where it is declared. The process of how names are resolved to entities is described in the <u>name resolution</u> chapter. More information on "drop scopes" used for the purpose of running destructors may be found in the <u>destructors</u> chapter.

r[names.scopes.items]

## **Item scopes**

r[names.scopes.items.module] The name of an <u>item</u> declared directly in a <u>module</u> has a scope that extends from the start of the module to the end of the module. These items are also members of the module and can be referred to with a <u>path</u> leading from their module.

r[names.scopes.items.statement] The name of an item declared as a <u>statement</u> has a scope that extends from the start of the block the item statement is in until the end of the block.

r[names.scopes.items.duplicate] It is an error to introduce an item with a duplicate name of another item in the same <u>namespace</u> within the same module or block. <u>Asterisk glob imports</u> have special behavior for dealing with duplicate names and shadowing, see the linked chapter for more details.

r[names.scopes.items.shadow-prelude] Items in a module may shadow items in a <u>prelude</u>.

r[names.scopes.items.nested-modules] Item names from outer modules are not in scope within a nested module. A <u>path</u> may be used to refer to an item in another module.

r[names.scopes.associated-items]

## **Associated item scopes**

r[names.scopes.associated-items.scope] <u>Associated items</u> are not scoped and can only be referred to by using a <u>path</u> leading from the type or trait they are associated with. <u>Methods</u> can also be referred to via <u>call</u> <u>expressions</u>.

r[names.scopes.associated-items.duplicate] Similar to items within a module or block, it is an error to introduce an item within a trait or implementation that is a duplicate of another item in the trait or impl in the same namespace.

r[names.scopes.pattern-bindings]

# **Pattern binding scopes**

The scope of a local variable <u>pattern</u> binding depends on where it is used:

r[names.scopes.pattern-bindings.let]

- <u>let statement</u> bindings range from just after the <u>let</u> statement until the end of the block where it is declared. r[names.scopes.patternbindings.parameter]
- <u>Function parameter</u> bindings are within the body of the function. r[names.scopes.pattern-bindings.closure]
- <u>Closure parameter</u> bindings are within the closure body. r[names.scopes.pattern-bindings.loop]
- <u>for</u> bindings are within the loop body. r[names.scopes.patternbindings.let-chains]
- if let and while let bindings are valid in the following conditions as well as the consequent block. r[names.scopes.pattern-bindings.match-arm]
- <u>match</u> arms bindings are within the <u>match guard</u> and the match arm expression.

r[names.scopes.pattern-bindings.items] Local variable scopes do not extend into item declarations.

# Pattern binding shadowing

r[names.scopes.pattern-bindings.shadow] Pattern bindings are allowed to shadow any name in scope with the following exceptions which are an error:

- <u>Const generic parameters</u>
- <u>Static items</u>
- <u>Const items</u>
- Constructors for <u>structs</u> and <u>enums</u>

The following example illustrates how local bindings can shadow item declarations:

```
fn shadow_example() {
```

// Since there are no local variables in scope yet, this
resolves to the function.

```
foo(); // prints `function`
```

```
let foo = || println!("closure");
```

```
fn foo() { println!("function"); }
```

// This resolves to the local closure since it shadows the
item.

```
foo(); // prints `closure`
```

}

r[names.scopes.generic-parameters]

## **Generic parameter scopes**

r[names.scopes.generic-parameters.param-list] Generic parameters are declared in a [GenericParams] list. The scope of a generic parameter is within the item it is declared on.

r[names.scopes.generic-parameters.order-independent] All parameters are in scope within the generic parameter list regardless of the order they are declared. The following shows some examples where a parameter may be referenced before it is declared:

```
// The 'b bound is referenced before it is declared.
fn params_scope<'a: 'b, 'b>() {}
# trait SomeTrait<const Z: usize> {}
// The const N is referenced in the trait bound before it is
declared.
fn f<T: SomeTrait<N>, const N: usize>() {}
```

r[names.scopes.generic-parameters.bounds] Generic parameters are also in scope for type bounds and where clauses, for example:

```
# trait SomeTrait<'a, T> {}
// The <'a, U> for `SomeTrait` refer to the 'a and U
parameters of `bounds_scope`.
fn bounds_scope<'a, T: SomeTrait<'a, U>, U>() {}
fn where_scope<'a, T, U>()
where T: SomeTrait<'a, U>
{}
```

r[names.scopes.generic-parameters.inner-items] It is an error for <u>items</u> declared inside a function to refer to a generic parameter from their outer scope.

```
fn example<T>() {
    fn inner(x: T) {} // ERROR: can't use generic parameters
from outer function
}
```

## **Generic parameter shadowing**

r[names.scopes.generic-parameters.shadow] It is an error to shadow a generic parameter with the exception that items declared within functions are allowed to shadow generic parameter names from the function.

```
fn example<'a, T, const N: usize>() {
      // Items within functions are allowed to shadow generic
 parameter in scope.
     fn inner_lifetime<'a>() {} // OK
     fn inner_type<T>() {} // OK
     fn inner_const<const N: usize>() {} // OK
 }
trait SomeTrait<'a, T, const N: usize> {
    fn example_lifetime<'a>() {} // ERROR: 'a is already in use
    fn example_type<T>() {} // ERROR: T is already in use
      fn example_const<const N: usize>() {} // ERROR:
                                                          N is
alreadv in use
      fn example mixed<const T: usize>() {} // ERROR: T is
already in use
}
```

```
r[names.scopes.lifetimes]
```

## Lifetime scopes

Lifetime parameters are declared in a [GenericParams] list and <u>higher-</u> <u>ranked trait bounds</u>.

r[names.scopes.lifetimes.special] The 'static lifetime and placeholder
lifetime '\_ have a special meaning and cannot be declared as a parameter.

#### Lifetime generic parameter scopes

r[names.scopes.lifetimes.generic] <u>Constant</u> and <u>static</u> items and <u>const</u> <u>contexts</u> only ever allow 'static lifetime references, so no other lifetime may be in scope within them. <u>Associated consts</u> do allow referring to lifetimes declared in their trait or implementation.

#### Higher-ranked trait bound scopes

r[names.scopes.lifetimes.higher-ranked] The scope of a lifetime parameter declared as a <u>higher-ranked trait bound</u> depends on the scenario where it is used.

- As a [TypeBoundWhereClauseItem] the declared lifetimes are in scope in the type and the type bounds.
- As a [TraitBound] the declared lifetimes are in scope within the bound type path.
- As a [BareFunctionType] the declared lifetimes are in scope within the function parameters and return type.

```
# trait Trait<'a>{}
fn where_clause<T>()
    // 'a is in scope in both the type and the type bounds.
    where for <'a> &'a T: Trait<'a>
{}
fn bound<T>()
    // 'a is in scope within the bound.
    where T: for <'a> Trait<'a>
{}
# struct Example<'a> {
    field: &'a u32
# }
// 'a is in scope in both the parameters and return type.
type FnExample = for<'a> fn(x: Example<'a>) -> Example<'a>;
```

### **Impl trait restrictions**

r[names.scopes.lifetimes.impl-trait] <u>Impl trait</u> types can only reference lifetimes declared on a function or implementation.

```
# trait Trait1 {
# type Item;
# }
```

```
# trait Trait2<'a> {}
#
# struct Example;
#
# impl Trait1 for Example {
      type Item = Element;
#
# }
#
# struct Element;
# impl<'a> Trait2<'a> for Element {}
#
// The `impl Trait2` here is not allowed to refer to 'b but it
is allowed to
// refer to 'a.
fn foo<'a>() -> impl for<'b> Trait1<Item = impl Trait2<'a> +
use<'a>> {
   // ...
     Example
#
}
```

r[names.scopes.loop-label]
## Loop label scopes

r[names.scopes.loop-label.scope] <u>Loop labels</u> may be declared by a <u>loop</u> <u>expression</u>. The scope of a loop label is from the point it is declared till the end of the loop expression. The scope does not extend into <u>items</u>, <u>closures</u>, <u>async blocks</u>, <u>const arguments</u>, <u>const contexts</u>, and the iterator expression of the defining <u>for loop</u>.

```
'a: for n in 0..3 {
    if n % 2 == 0 {
        break 'a;
    }
    fn inner() {
        // Using 'a here would be an error.
        // break 'a;
    }
}
// The label is in scope for the expression of `while` loops.
'a: while break 'a {}
                             // Loop does not run.
'a: while let _ = break 'a {} // Loop does not run.
// The label is not in scope in the defining `for` loop:
'a: for outer in 0..5 {
    // This will break the outer loop, skipping the inner loop
and stopping
    // the outer loop.
    'a: for inner in { break 'a; 0..1 } {
        println!("{}", inner); // This does not run.
    }
    println!("{}", outer); // This does not run, either.
}
```

r[names.scopes.loop-label.shadow] Loop labels may shadow labels of the same name in outer scopes. References to a label refer to the closest definition.

```
// Loop label shadowing example.
'a: for outer in 0..5 {
    'a: for inner in 0..5 {
        // This terminates the inner loop, but the outer loop
continues to run.
        break 'a;
    }
}
```

r[names.scopes.prelude]

### **Prelude scopes**

r[names.scopes.prelude.intro] <u>Preludes</u> bring entities into scope of every module. The entities are not members of the module, but are implicitly queried during <u>name resolution</u>.

r[names.scopes.prelude.shadow] The prelude names may be shadowed by declarations in a module.

r[names.scopes.prelude.layers] The preludes are layered such that one shadows another if they contain entities of the same name. The order that preludes may shadow other preludes is the following where earlier entries may shadow later ones:

- 1. Extern prelude
- 2. <u>Tool prelude</u>
- 3. <u>macro use prelude</u>
- 4. Standard library prelude
- 5. <u>Language prelude</u>
- r[names.scopes.macro\_rules]

#### macro\_rules scopes

The scope of macro\_rules macros is described in the <u>Macros By</u> <u>Example</u> chapter. The behavior depends on the use of the <u>macro use</u> and <u>macro export</u> attributes.

r[names.scopes.derive]

### **Derive macro helper attributes**

r[names.scopes.derive.scope] <u>Derive macro helper attributes</u> are in scope in the item where their corresponding <u>derive attribute</u> is specified. The scope extends from just after the <u>derive</u> attribute to the end of the item.

r[names.scopes.derive.shadow] Helper attributes shadow other attributes of the same name in scope.

r[names.scopes.self]

#### Self scope

r[names.scopes.self.intro] Although <u>Self</u> is a keyword with special meaning, it interacts with name resolution in a way similar to normal names.

r[names.scopes.self.def-scope] The implicit Self type in the definition of a <u>struct</u>, <u>enum</u>, <u>union</u>, <u>trait</u>, or <u>implementation</u> is treated similarly to a <u>generic parameter</u>, and is in scope in the same way as a generic type parameter.

r[names.scopes.self.impl-scope] The implicit Self constructor in the value <u>namespace</u> of an <u>implementation</u> is in scope within the body of the implementation (the implementation's <u>associated items</u>).

```
// Self type within struct definition.
struct Recursive {
   f1: Option<Box<Self>>
}
// Self type within generic parameters.
struct SelfGeneric<T: Into<Self>>(T);
// Self value constructor within an implementation.
struct ImplExample();
impl ImplExample {
   fn example() -> Self { // Self type
      Self() // Self value constructor
   }
}
```

r[names.preludes]

# Preludes

r[names.preludes.intro] A *prelude* is a collection of names that are automatically brought into scope of every module in a crate.

These prelude names are not part of the module itself: they are implicitly queried during <u>name resolution</u>. For example, even though something like [Box] is in scope in every module, you cannot refer to it as self::Box because it is not a member of the current module.

r[names.preludes.kinds] There are several different preludes:

- <u>Standard library prelude</u>
- Extern prelude
- <u>Language prelude</u>
- <u>macro use</u> <u>prelude</u>
- <u>Tool prelude</u>

r[names.preludes.std]

# **Standard library prelude**

r[names.preludes.std.intro] Each crate has a standard library prelude, which consists of the names from a single standard library module.

r[names.preludes.std.module] The module used depends on the crate's edition, and on whether the <u>no std attribute</u> is applied to the crate:

Edition	no_std not applied	no_std applied
2015	[std::prelude::rust_ 2015]	[core::prelude::rust_ 2015]
2018	[std::prelude::rust_ 2018]	[core::prelude::rust_ 2018]
2021	[std::prelude::rust_ 2021]	[core::prelude::rust_ 2021]
2024	[std::prelude::rust_ 2024]	[core::prelude::rust_ 2024]

[!NOTE] [std::prelude::rust\_2015] and [std::prelude::rust\_2018] have the same contents as [std::prelude::v1].

[core::prelude::rust\_2015] and [core::prelude::rust\_2018]
have the same contents as [core::prelude::v1].

r[names.preludes.extern]

## **Extern prelude**

r[names.preludes.extern.intro] External crates imported with <u>extern</u> <u>crate</u> in the root module or provided to the compiler (as with the --<u>extern</u> flag with <u>rustc</u>) are added to the *extern prelude*. If imported with an alias such as <u>extern crate orig\_name</u> as <u>new\_name</u>, then the symbol <u>new\_name</u> is instead added to the prelude.

r[names.preludes.extern.core] The [core] crate is always added to the extern prelude.

r[names.preludes.extern.std] The [std] crate is added as long as the <u>no std attribute</u> is not specified in the crate root.

r[names.preludes.extern.edition2018]

[!EDITION-2018] In the 2015 edition, crates in the extern prelude cannot be referenced via <u>use declarations</u>, so it is generally standard practice to include <u>extern crate</u> declarations to bring them into scope.

Beginning in the 2018 edition, <u>use declarations</u> can reference crates in the extern prelude, so it is considered unidiomatic to use <u>extern</u> <u>crate</u>.

[!NOTE] Additional crates that ship with rustc, such as [alloc], and <u>test</u>, are not automatically included with the --extern flag when using Cargo. They must be brought into scope with an extern crate declaration, even in the 2018 edition.

```
extern crate alloc;
use alloc::rc::Rc;
```

Cargo does bring in proc\_macro to the extern prelude for procmacro crates only.

r[names.preludes.extern.no\_std]

### The no\_std attribute

r[names.preludes.extern.no\_std.intro] By default, the standard library is automatically included in the crate root module. The [std] crate is added to the root, along with an implicit <u>macro use attribute</u> pulling in all macros exported from std into the <u>macro use prelude</u>. Both [core] and [std] are added to the <u>extern prelude</u>.

r[names.preludes.extern.no\_std.allowed-positions] The *no\_std* <u>attribute</u> may be applied at the crate level to prevent the [std] crate from being automatically added into scope.

It does three things:

r[names.preludes.extern.no\_std.extern]

- Prevents std from being added to the <u>extern prelude</u>. r[names.preludes.extern.no\_std.module]
- Affects which module is used to make up the <u>standard library prelude</u> (as described above). r[names.preludes.extern.no\_std.core]
- Injects the [core] crate into the crate root instead of [std], and pulls in all macros exported from core in the <u>macro use prelude</u>.

[!NOTE] Using the core prelude over the standard prelude is useful when either the crate is targeting a platform that does not support the standard library or is purposefully not using the capabilities of the standard library. Those capabilities are mainly dynamic memory allocation (e.g. Box and Vec) and file and network capabilities (e.g. std::fs and std::io).

[!WARNING] Using no\_std does not prevent the standard library from being linked in. It is still valid to put extern crate std; into the crate and dependencies can also link it in.

r[names.preludes.lang]

## Language prelude

r[names.preludes.lang.intro] The language prelude includes names of types and attributes that are built-in to the language. The language prelude is always in scope.

r[names.preludes.lang.entities] It includes the following:

• <u>Type namespace</u>

- <u>Boolean type</u> --- bool
- <u>Textual types</u> --- char and str
- Integer types --- i8, i16, i32, i64, i128, u8, u16, u32, u64, u128
- Machine-dependent integer types --- usize and isize
- <u>floating-point types</u> --- f32 and f64
- <u>Macro namespace</u>

• <u>Built-in attributes</u>

r[names.preludes.macro\_use]

### macro\_use prelude

r[names.preludes.macro\_use.intro] The macro\_use prelude includes macros from external crates that were imported by the macro\_use attribute applied to an extern crate.

r[names.preludes.tool]

## **Tool prelude**

r[names.preludes.tool.intro] The tool prelude includes tool names for external tools in the <u>type namespace</u>. See the <u>tool attributes</u> section for more details.

r[names.preludes.no\_implicit\_prelude]

### The no\_implicit\_prelude attribute

r[names.preludes.no\_implicit\_prelude.intro] The *no\_implicit\_prelude attribute* may be applied at the crate level or on a module to indicate that it should not automatically bring the <u>standard library prelude</u>, <u>extern prelude</u>, or <u>tool prelude</u> into scope for that module or any of its descendants.

r[names.preludes.no\_implicit\_prelude.lang] This attribute does not affect the <u>language prelude</u>.

r[names.preludes.no\_implicit\_prelude.edition2018]

[!EDITION-2018] In the 2015 edition, the no\_implicit\_prelude attribute does not affect the <u>macro use\_prelude</u>, and all macros exported from the standard library are still included in the <u>macro\_use</u> prelude. Starting in the 2018 edition, it will remove the <u>macro\_use</u> prelude.

r[paths]

# **Paths**

r[paths.intro] A *path* is a sequence of one or more path segments separated by :: tokens. Paths are used to refer to <u>items</u>, values, <u>types</u>, <u>macros</u>, and <u>attributes</u>.

Two examples of simple paths consisting of only identifier segments: x; x::y::z;

# **Types of paths**

r[paths.simple]

# **Simple Paths**

```
r[paths.simple.syntax]
SimplePath ->
`::`? SimplePathSegment (`::` SimplePathSegment)*
```

```
SimplePathSegment ->
```

```
IDENTIFIER | `super` | `self` | `crate` | `$crate`
```

r[paths.simple.intro] Simple paths are used in <u>visibility</u> markers, <u>attributes</u>, <u>macros</u>, and <u>use</u> items. For example:

```
use std::io::{self, Write};
mod m {
    #[clippy::cyclomatic_complexity = "0"]
    pub (in super) fn f1() {}
}
```

r[paths.expr]

# **Paths in expressions**

```
r[paths.expr.syntax]
PathInExpression ->
   `::`? PathExprSegment (`::` PathExprSegment)*
PathExprSegment ->
   PathIdentSegment (`::` GenericArgs)?
PathIdentSegment ->
   IDENTIFIER | `super` | `self` | `Self` | `crate` | `$crate`
GenericArgs ->
        `<` `>`
        | `<` ( GenericArg `,` )* GenericArg `,`? `>`
```

GenericArg ->

Lifetime | Type | GenericArgsConst | GenericArgsBinding | GenericArgsBounds

```
GenericArgsConst ->
```

BlockExpression

| LiteralExpression

| `-` LiteralExpression

| SimplePathSegment

```
GenericArgsBinding ->
```

```
IDENTIFIER GenericArgs? `=` Type
```

GenericArgsBounds ->

IDENTIFIER GenericArgs? `:` TypeParamBounds

r[paths.expr.intro] Paths in expressions allow for paths with generic arguments to be specified. They are used in various places in <u>expressions</u> and <u>patterns</u>.

r[paths.expr.turbofish] The :: token is required before the opening < for generic arguments to avoid ambiguity with the less-than operator. This is colloquially known as "turbofish" syntax.

```
(0..10).collect::<Vec<_>>();
Vec::<u8>::with_capacity(1024);
```

r[paths.expr.argument-order] The order of generic arguments is restricted to lifetime arguments, then type arguments, then const arguments, then equality constraints.

r[paths.expr.complex-const-params] Const arguments must be surrounded by braces unless they are a <u>literal</u>, an <u>inferred const</u>, or a single segment path. An <u>inferred const</u> may not be surrounded by braces.

```
mod m {
    pub const C: usize = 1;
}
const C: usize = m::C;
```

[!NOTE] In a generic argument list, an <u>inferred const</u> is parsed as an [inferred type][InferredType] but then semantically treated as a separate kind of <u>const generic argument</u>.

r[paths.expr.impl-trait-params] The synthetic type parameters corresponding to impl Trait types are implicit, and these cannot be explicitly specified.

r[paths.qualified]

# **Qualified** paths

```
r[paths.qualified.syntax]
QualifiedPathInExpression -> QualifiedPathType (`::`
PathExprSegment)+
QualifiedPathType -> `<` Type (`as` TypePath)? `>`
QualifiedPathInType -> QualifiedPathType (`::`
TypePathSegment)+
```

r[paths.qualified.intro] Fully qualified paths allow for disambiguating the path for <u>trait implementations</u> and for specifying <u>canonical paths</u>. When used in a type specification, it supports using the type syntax specified below.

```
struct S;
impl S {
    fn f() { println!("S"); }
}
trait T1 {
    fn f() { println!("T1 f"); }
}
impl T1 for S {}
trait T2 {
    fn f() { println!("T2 f"); }
}
impl T2 for S {}
S::f(); // Calls the inherent impl.
<S as T1>::f(); // Calls the T1 trait function.
<S as T2>::f(); // Calls the T2 trait function.
```

r[paths.type]

# Paths in types

r[paths.type.syntax]

```
TypePath -> `::`? TypePathSegment (`::` TypePathSegment)*
```

```
TypePathSegment -> PathIdentSegment (`::`? (GenericArgs |
TypePathFn))?
```

```
TypePathFn -> `(` TypePathFnInputs? `)` (`->` TypeNoBounds)?
```

```
TypePathFnInputs -> Type (`,` Type)* `,`?
```

r[paths.type.intro] Type paths are used within type definitions, trait bounds, type parameter bounds, and qualified paths.

r[paths.type.turbofish] Although the :: token is allowed before the generics arguments, it is not required because there is no ambiguity like there is in [PathInExpression].

```
# mod ops {
#
      pub struct Range<T> {f1: T}
      pub trait Index<T> {}
#
      pub struct Example<'a> {f1: &'a i32}
#
# }
# struct S;
impl ops::Index<ops::Range<usize>> for S { /*...*/ }
fn i<'a>() -> impl Iterator<Item = ops::Example<'a>> {
    // ...
     const EXAMPLE: Vec<ops::Example<'static>> = Vec::new();
#
     EXAMPLE.into_iter()
#
}
type G = std::boxed::Box<dyn std::ops::FnOnce(isize)</pre>
                                                              ->
isize>;
```

r[paths.qualifiers]

## **Path qualifiers**

Paths can be denoted with various leading qualifiers to change the meaning of how it is resolved.

r[paths.qualifiers.global-root]

::

r[paths.qualifiers.global-root.intro] Paths starting with :: are considered to be *global paths* where the segments of the path start being resolved from a place which differs based on edition. Each identifier in the path must resolve to an item.

r[paths.qualifiers.global-root.edition2018]

[!EDITION-2018] In the 2015 Edition, identifiers resolve from the "crate root" (crate:: in the 2018 edition), which contains a variety of different items, including external crates, default crates such as std or core, and items in the top level of the crate (including use imports).

Beginning with the 2018 Edition, paths starting with **::** resolve from crates in the <u>extern prelude</u>. That is, they must be followed by the name of a crate.

```
pub fn foo() {
    // In the 2018 edition, this accesses `std` via the extern
prelude.
    // In the 2015 edition, this accesses `std` via the crate
root.
    let now = ::std::time::Instant::now();
    println!("{:?}", now);
}
// 2015 Edition
mod a {
    pub fn foo() {}
}
mod b {
    pub fn foo() {
```

#### self

r[paths.qualifiers.mod-self.intro] self resolves the path relative to the current module.

r[paths.qualifiers.mod-self.restriction] self can only be used as the first segment, without a preceding ::.

r[paths.qualifiers.self-pat] In a method body, a path which consists of a single self segment resolves to the method's self parameter.

```
fn foo() {}
fn bar() {
    self::foo();
}
struct S(bool);
impl S {
    fn baz(self) {
        self.0;
    }
# fn main() {}
```

r[paths.qualifiers.type-self]

### Self

r[paths.qualifiers.type-self.intro] Self, with a capital "S", is used to refer to the current type being implemented or defined. It may be used in the following situations:

r[paths.qualifiers.type-self.trait]

- In a <u>trait</u> definition, it refers to the type implementing the trait. r[paths.qualifiers.type-self.impl]
- In an <u>implementation</u>, it refers to the type being implemented. When implementing a tuple or unit <u>struct</u>, it also refers to the constructor in the <u>value namespace</u>.

r[paths.qualifiers.type-self.type]

• In the definition of a <u>struct</u>, <u>enumeration</u>, or <u>union</u>, it refers to the type being defined. The definition is not allowed to be infinitely recursive (there must be an indirection).

r[paths.qualifiers.type-self.scope] The scope of Self behaves similarly to a generic parameter; see the <u>Self scope</u> section for more details.

r[paths.qualifiers.type-self.allowed-positions] Self can only be used as the first segment, without a preceding :::.

r[paths.qualifiers.type-self.no-generics] The Self path cannot include
generic arguments (as in Self::<i32>).

```
trait T {
   type Item;
   const C: i32;
   // `Self` will be whatever type that implements `T`.
   fn new() -> Self;
        // `Self::Item` will be the type alias in the
implementation.
   fn f(&self) -> Self::Item;
}
struct S;
impl T for S {
   type Item = i32;
   const C: i32 = 9;
   fn new() -> Self { // `Self` is the type `S`.
       S
   }
    fn f(&self) -> Self::Item { // `Self::Item` is the type
```

```
`i32`.
                                  // `Self::C` is the constant
        Self::C
value `9`.
    }
}
// `Self` is in scope within the generics of a trait
definition,
// to refer to the type being defined.
trait Add<Rhs = Self> {
    type Output;
    // `Self` can also reference associated items of the
    // type being implemented.
    fn add(self, rhs: Rhs) -> Self::Output;
}
struct NonEmptyList<T> {
    head: T,
    // A struct can reference itself (as long as it is not
    // infinitely recursive).
    tail: Option<Box<Self>>,
}
```

r[paths.qualifiers.super]

#### super

r[paths.qualifiers.super.intro] super in a path resolves to the parent
module.

r[paths.qualifiers.super.allowed-positions] It may only be used in leading segments of the path, possibly after an initial self segment.

```
mod a {
    pub fn foo() {}
}
mod b {
    pub fn foo() {
```

```
super::a::foo(); // call a's foo function
}
# fn main() {}
```

r[paths.qualifiers.super.repetition] super may be repeated several times after the first super or self to refer to ancestor modules.

r[paths.qualifiers.crate]

#### crate

r[paths.qualifiers.crate.intro] crate resolves the path relative to the current crate.

r[paths.qualifiers.crate.allowed-positions] crate can only be used as the first segment, without a preceding ::.

```
fn foo() {}
mod a {
    fn bar() {
        crate::foo();
    }
}
# fn main() {}
```

r[paths.qualifiers.macro-crate]

#### \$crate

r[paths.qualifiers.macro-crate.allowed-positions] <u>\$crate</u> is only used within <u>macro transcribers</u>, and can only be used as the first segment, without a preceding ::.

r[paths.qualifiers.macro-crate.hygiene] <u>\$crate</u> will expand to a path to access items from the top level of the crate where the macro is defined, regardless of which crate the macro is invoked.

```
pub fn increment(x: u32) -> u32 {
    x + 1
}
#[macro_export]
macro_rules! inc {
    ($x:expr) => ( $crate::increment($x) )
}
# fn main() { }
```

r[paths.canonical]

## **Canonical paths**

r[paths.canonical.intro] Items defined in a module or implementation have a *canonical path* that corresponds to where within its crate it is defined.

r[paths.canonical.alias] All other paths to these items are aliases.

r[paths.canonical.def] The canonical path is defined as a *path prefix* appended by the path segment the item itself defines.

r[paths.canonical.non-canonical] Implementations and use declarations do not have canonical paths, although the items that implementations define do have them. Items defined in block expressions do not have canonical paths. Items defined in a module that does not have a canonical path do not have a canonical path. Associated items defined in an implementation that refers to an item without a canonical path, e.g. as the implementing type, the trait being implemented, a type parameter or bound on a type parameter, do not have canonical paths.

r[paths.canonical.module-prefix] The path prefix for modules is the canonical path to that module.

r[paths.canonical.bare-impl-prefix] For bare implementations, it is the canonical path of the item being implemented surrounded by angle (<>) brackets.

r[paths.canonical.trait-impl-prefix] For <u>trait implementations</u>, it is the canonical path of the item being implemented followed by as followed by the canonical path to the trait all surrounded in angle (<>) brackets.

r[paths.canonical.local-canonical-path] The canonical path is only meaningful within a given crate. There is no global namespace across crates; an item's canonical path merely identifies it within the crate.

```
// Comments show the canonical path of the item.
mod a { // crate::a
   pub struct Struct; // crate::a::Struct
   pub trait Trait { // crate::a::Trait
```

```
fn f(&self); // crate::a::Trait::f
    }
    impl Trait for Struct {
                  fn f(&self) {} // <crate::a::Struct</pre>
                                                             as
crate::a::Trait>::f
    }
    impl Struct {
        fn g(&self) {} // <crate::a::Struct>::g
    }
}
mod without { // crate::without
    fn canonicals() { // crate::without::canonicals
        struct OtherStruct; // None
        trait OtherTrait { // None
            fn g(&self); // None
        }
        impl OtherTrait for OtherStruct {
            fn g(&self) {} // None
        }
        impl OtherTrait for crate::a::Struct {
            fn g(&self) {} // None
        }
        impl crate::a::Trait for OtherStruct {
            fn f(&self) {} // None
        }
    }
}
# fn main() {}
```

# Name resolution

[!NOTE] This is a placeholder for future expansion.

r[vis]

# **Visibility and Privacy**

```
r[vis.syntax]
Visibility ->
```

r[vis.intro] These two terms are often used interchangeably, and what they are attempting to convey is the answer to the question "Can this item be used at this location?"

r[vis.name-hierarchy] Rust's name resolution operates on a global hierarchy of namespaces. Each level in the hierarchy can be thought of as some item. The items are one of those mentioned above, but also include external crates. Declaring or defining a new module can be thought of as inserting a new tree into the hierarchy at the location of the definition.

r[vis.privacy] To control whether interfaces can be used across modules, Rust checks each use of an item to see whether it should be allowed or not. This is where privacy warnings are generated, or otherwise "you used a private item of another module and weren't allowed to."

r[vis.default] By default, everything is *private*, with two exceptions: Associated items in a pub Trait are public by default; Enum variants in a pub enum are also public by default. When an item is declared as pub, it can be thought of as being accessible to the outside world. For example:

```
# fn main() {}
// Declare a private struct
struct Foo;
// Declare a public struct with a private field
pub struct Bar {
    field: i32,
}
```

```
// Declare a public enum with two public variants
pub enum State {
    PubliclyAccessibleState,
    PubliclyAccessibleState2,
}
```

r[vis.access] With the notion of an item being either public or private, Rust allows item accesses in two cases:

- 1. If an item is public, then it can be accessed externally from some module m if you can access all the item's ancestor modules from m. You can also potentially be able to name the item through re-exports. See below.
- 2. If an item is private, it may be accessed by the current module and its descendants.

These two cases are surprisingly powerful for creating module hierarchies exposing public APIs while hiding internal implementation details. To help explain, here's a few use cases and what they would entail:

- A library developer needs to expose functionality to crates which link against their library. As a consequence of the first case, this means that anything which is usable externally must be pub from the root down to the destination item. Any private item in the chain will disallow external accesses.
- A crate needs a global available "helper module" to itself, but it doesn't want to expose the helper module as a public API. To accomplish this, the root of the crate's hierarchy would have a private module which then internally has a "public API". Because the entire crate is a descendant of the root, then the entire local crate can access this private module through the second case.
- When writing unit tests for a module, it's often a common idiom to have an immediate child of the module to-be-tested named mod test. This module could access any items of the parent module through the second case, meaning that internal implementation details could also be seamlessly tested from the child module.

In the second case, it mentions that a private item "can be accessed" by the current module and its descendants, but the exact meaning of accessing an item depends on what the item is.

r[vis.usage] Accessing a module, for example, would mean looking inside of it (to import more items). On the other hand, accessing a function would mean that it is invoked. Additionally, path expressions and import statements are considered to access an item in the sense that the import/expression is only valid if the destination is in the current visibility scope.

Here's an example of a program which exemplifies the three cases outlined above:

```
// This module is private, meaning that no external crate can
access this
// module. Because it is private at the root of this current
crate, however, any
// module in the crate may access any publicly visible item in
this module.
mod crate_helper_module {
     // This function can be used by anything in the current
crate
    pub fn crate_helper() {}
    // This function *cannot* be used by anything else in the
crate. It is not
    // publicly visible outside of the `crate_helper_module`,
so only this
    // current module and its descendants may access it.
    fn implementation_detail() {}
}
// This function is "public to the root" meaning that it's
available to external
// crates linking against this one.
```

```
pub fn public_api() {}
```
```
// Similarly to 'public_api', this module is public so
external crates may look
// inside of it.
pub mod submodule {
    use crate::crate_helper_module;
    pub fn my_method() {
        // Any item in the local crate may invoke the helper
module's public
         // interface through a combination of the two rules
above.
        crate_helper_module::crate_helper();
    }
    // This function is hidden to any module which is not a
descendant of
    // `submodule`
    fn my_implementation() {}
    #[cfg(test)]
    mod test {
        #[test]
        fn test my implementation() {
                 // Because this module is a descendant of
`submodule`, it's allowed
             // to access private items inside of `submodule`
without a privacy
            // violation.
            super::my_implementation();
        }
    }
}
# fn main() {}
```

For a Rust program to pass the privacy checking pass, all paths must be valid accesses given the two rules above. This includes all use statements, expressions, types, etc.

r[vis.scoped]

### 

r[vis.scoped.intro] In addition to public and private, Rust allows users to declare an item as visible only within a given scope. The rules for pub restrictions are as follows:

r[vis.scoped.in]

• pub(in path) makes an item visible within the provided path. path must be a simple path which resolves to an ancestor module of the item whose visibility is being declared. Each identifier in path must refer directly to a module (not to a name introduced by a use statement).

r[vis.scoped.crate]

- pub(crate) makes an item visible within the current crate.
   r[vis.scoped.super]
- pub(super) makes an item visible to the parent module. This is equivalent to pub(in super).
   r[vis.scoped.self]
- pub(self) makes an item visible to the current module. This is equivalent to pub(in self) or not using pub at all.
   r[vis.scoped.edition2018]

```
[!EDITION-2018] Starting with the 2018 edition, paths for pub(in path) must start with crate, self, or super. The 2015 edition may also use paths starting with :: or modules from the crate root.
```

```
Here's an example:
pub mod outer_mod {
    pub mod inner_mod {
        // This function is visible within `outer_mod`
```

```
pub(in crate::outer_mod) fn outer_mod_visible_fn() {}
           // Same as above, this is only valid in the 2015
edition.
        pub(in outer_mod) fn outer_mod_visible_fn_2015() {}
        // This function is visible to the entire crate
        pub(crate) fn crate visible fn() {}
        // This function is visible within `outer mod`
        pub(super) fn super_mod_visible_fn() {
            // This function is visible since we're in the same
`mod`
            inner_mod_visible_fn();
        }
        // This function is visible only within `inner_mod`,
        // which is the same as leaving it private.
        pub(self) fn inner_mod_visible_fn() {}
    }
    pub fn foo() {
        inner mod::outer mod visible fn();
        inner mod::crate visible fn();
        inner_mod::super_mod_visible_fn();
           // This function is no longer visible since we're
outside of `inner mod`
        // Error! `inner mod visible fn` is private
        //inner_mod::inner_mod_visible_fn();
    }
}
fn bar() {
    // This function is still visible since we're in the same
crate
    outer_mod::inner_mod::crate_visible_fn();
```

// This function is no longer visible since we're outside
of `outer\_mod`

```
// Error! `super_mod_visible_fn` is private
//outer_mod::inner_mod::super_mod_visible_fn();
```

// This function is no longer visible since we're outside
of `outer\_mod`

```
// Error! `outer_mod_visible_fn` is private
//outer_mod::inner_mod::outer_mod_visible_fn();
```

```
outer_mod::foo();
```

```
}
```

```
fn main() { bar() }
```

[!NOTE] This syntax only adds another restriction to the visibility of an item. It does not guarantee that the item is visible within all parts of the specified scope. To access an item, all of its parent items up to the current scope must still be visible as well.

```
r[vis.reexports]
```

#### **Re-exporting and Visibility**

r[vis.reexports.intro] Rust allows publicly re-exporting items through a pub use directive. Because this is a public directive, this allows the item to be used in the current module through the rules above. It essentially allows public access into the re-exported item. For example, this program is valid:

```
pub use self::implementation::api;
mod implementation {
    pub mod api {
        pub fn f() {}
    }
}
# fn main() {}
```

This means that any external crate referencing implementation::api::f would receive a privacy violation, while the path api::f would be allowed.

r[vis.reexports.private-item] When re-exporting a private item, it can be thought of as allowing the "privacy chain" being short-circuited through the reexport instead of passing through the namespace hierarchy as it normally would. r[memory]

## **Memory model**

[!WARNING] The memory model of Rust is incomplete and not fully decided.

r[memory.bytes]

#### **Bytes**

r[memory.bytes.intro] The most basic unit of memory in Rust is a byte.

[!NOTE] While bytes are typically lowered to hardware bytes, Rust uses an "abstract" notion of bytes that can make distinctions which are absent in hardware, such as being uninitialized, or storing part of a pointer. Those distinctions can affect whether your program has undefined behavior, so they still have tangible impact on how compiled Rust programs behave.

r[memory.bytes.contents] Each byte may have one of the following values:

r[memory.bytes.init]

• An initialized byte containing a u8 value and optional [provenance] [std::ptr#provenance],

r[memory.bytes.uninit]

• An uninitialized byte.

[!NOTE] The above list is not yet guaranteed to be exhaustive.

r[alloc]

## **Memory allocation and lifetime**

r[alloc.static] The *items* of a program are those functions, modules, and types that have their value calculated at compile-time and stored uniquely in the memory image of the rust process. Items are neither dynamically allocated nor freed.

r[alloc.dynamic] The *heap* is a general term that describes boxes. The lifetime of an allocation in the heap depends on the lifetime of the box values pointing to it. Since box values may themselves be passed in and out of frames, or stored in the heap, heap allocations may outlive the frame they are allocated within. An allocation in the heap is guaranteed to reside at a single location in the heap for the whole lifetime of the allocation - it will never be relocated as a result of moving a box value.

r[variable]

### Variables

r[variable.intro] A *variable* is a component of a stack frame, either a named function parameter, an anonymous <u>temporary</u>, or a named local variable.

r[variable.local] A *local variable* (or *stack-local* allocation) holds a value directly, allocated within the stack's memory. The value is a part of the stack frame.

r[variable.local-mut] Local variables are immutable unless declared otherwise. For example: let  $mut \times = ...$ .

r[variable.param-mut] Function parameters are immutable unless declared with mut. The mut keyword applies only to the following parameter. For example: [mut x, y] and fn f(mut x: Box<i32>, y: Box<i32>) declare one mutable variable x and one immutable variable y.

r[variable.init] Local variables are not initialized when allocated. Instead, the entire frame worth of local variables are allocated, on frameentry, in an uninitialized state. Subsequent statements within a function may or may not initialize the local variables. Local variables can be used only after they have been initialized through all reachable control flow paths.

In this next example, init\_after\_if is initialized after the <u>if</u> <u>expression</u> while <u>uninit\_after\_if</u> is not because it is not initialized in the else case.

```
# fn random_bool() -> bool { true }
fn initialization_example() {
    let init_after_if: ();
    let uninit_after_if: ();
    if random_bool() {
        init_after_if = ();
        uninit_after_if = ();
    } else {
        init_after_if = ();
    }
}
```

```
}
init_after_if; // ok
// uninit_after_if; // err: use of possibly uninitialized
`uninit_after_if`
}
```

r[panic]

### Panic

r[panic.intro] Rust provides a mechanism to prevent a function from returning normally, and instead "panic," which is a response to an error condition that is typically not expected to be recoverable within the context in which the error is encountered.

r[panic.lang-ops] Some language constructs, such as out-of-bounds <u>array</u> <u>indexing</u>, panic automatically.

r[panic.control] There are also language features that provide a level of control over panic behavior:

- A *panic handler* defines the behavior of a panic.
- **<u>FFI ABIs</u>** may alter how panics behave.

[!NOTE] The standard library provides the capability to explicitly panic via the [panic! macro][panic!].

r[panic.panic\_handler]

#### The panic\_handler attribute

r[panic.panic\_handler.intro] The *panic\_handler* attribute can be applied to a function to define the behavior of panics.

r[panic.panic\_handler.allowed-positions] The panic\_handler attribute can only be applied to a function with signature fn(&PanicInfo) -> !.

[!NOTE] The <u>PanicInfo</u> struct contains information about the location of the panic.

r[panic.panic\_handler.unique] There must be a single panic\_handler function in the dependency graph.

Below is shown a panic\_handler function that logs the panic message and then halts the thread.

#![no\_std]

```
use core::fmt::{self, Write};
use core::panic::PanicInfo;
struct Sink {
    // ..
    _0: (),
#
}
#
# impl Sink {
      fn new() -> Sink { Sink { _0: () }}
#
# }
#
# impl fmt::Write for Sink {
      fn write_str(&mut self, _: &str) -> fmt::Result { Ok(())
#
}
# }
#[panic_handler]
fn panic(info: &PanicInfo) -> ! {
```

```
let mut sink = Sink::new();
    // logs "panicked at '$reason', src/main.rs:27:4" to some
`sink`
    let _ = writeln!(sink, "{}", info);
    loop {}
}
```

r[panic.panic\_handler.std]

#### **Standard behavior**

r[panic.panic\_handler.std.kinds] std provides two different panic handlers:

- unwind --- unwinds the stack and is potentially recoverable.
- abort ---- aborts the process and is non-recoverable.

Not all targets may provide the unwind handler.

[!NOTE] The panic handler used when linking with std can be set with the <u>-C panic</u> CLI flag. The default for most targets is unwind.

The standard library's panic behavior can be modified at runtime with the [std::panic::set\_hook] function.

r[panic.panic\_handler.std.no\_std] Linking a <u>no\_std</u> binary, dylib, cdylib, or staticlib will require specifying your own panic handler.

r[panic.strategy]

#### **Panic strategy**

r[panic.strategy.intro] The *panic strategy* defines the kind of panic behavior that a crate is built to support.

[!NOTE] The panic strategy can be chosen in **rustc** with the <u>-c</u> panic CLI flag.

When generating a binary, dylib, cdylib, or staticlib and linking with std, the -C panic CLI flag also influences which <u>panic handler</u> is used.

[!NOTE] When compiling code with the abort panic strategy, the optimizer may assume that unwinding across Rust frames is impossible, which can result in both code-size and runtime speed improvements.

[!NOTE] See [link.unwinding] for restrictions on linking crates with different panic strategies. An implication is that crates built with the unwind strategy can use the abort panic handler, but the abort strategy cannot use the unwind panic handler.

r[panic.unwind]

#### Unwinding

r[panic.unwind.intro] Panicking may either be recoverable or non-recoverable, though it can be configured (by choosing a non-unwinding panic handler) to always be non-recoverable. (The converse is not true: the unwind handler does not guarantee that all panics are recoverable, only that panicking via the panic! macro and similar standard library mechanisms is recoverable.)

r[panic.unwind.destruction] When a panic occurs, the unwind handler "unwinds" Rust frames, just as C++'s throw unwinds C++ frames, until the panic reaches the point of recovery (for instance at a thread boundary). This means that as the panic traverses Rust frames, live objects in those frames that <u>implement Drop</u> will have their drop methods called. Thus, when normal execution resumes, no-longer-accessible objects will have been "cleaned up" just as if they had gone out of scope normally.

[!NOTE] As long as this guarantee of resource-cleanup is preserved, "unwinding" may be implemented without actually using the mechanism used by C++ for the target platform.

[!NOTE] The standard library provides two mechanisms for recovering from a panic, [std::panic::catch\_unwind] (which enables recovery within the panicking thread) and [std::thread::spawn] (which automatically sets up panic recovery for the spawned thread so that other threads may continue running).

r[panic.unwind.ffi]

#### **Unwinding across FFI boundaries**

r[panic.unwind.ffi.intro] It is possible to unwind across FFI boundaries using an <u>appropriate ABI declaration</u>. While useful in certain cases, this creates unique opportunities for undefined behavior, especially when multiple language runtimes are involved.

r[panic.unwind.ffi.undefined] Unwinding with the wrong ABI is undefined behavior:

- Causing an unwind into Rust code from a foreign function that was called via a function declaration or pointer declared with a non-unwinding ABI, such as "C", "system", etc. (For example, this case occurs when such a function written in C++ throws an exception that is uncaught and propagates to Rust.)
- Calling a Rust extern function that unwinds (with extern "Cunwind" or another ABI that permits unwinding) from code that does not support unwinding, such as code compiled with GCC or Clang using -fno-exceptions

r[panic.unwind.ffi.catch-foreign] Catching a foreign unwinding operation (such as a C++ exception) using [std::panic::catch\_unwind], [std::thread::JoinHandle::join], or by letting it propagate beyond the Rust main() function or thread root will have one of two behaviors, and it is unspecified which will occur:

- The process aborts.
- The function returns a [Result::Err] containing an opaque type.

[!NOTE] Rust code compiled or linked with a different instance of the Rust standard library counts as a "foreign exception" for the purpose of this guarantee. Thus, a library that uses panic! and is linked against one version of the Rust standard library, invoked from an application that uses a different version of the standard library, may cause the entire application to abort even if the library is only used within a child thread.

r[panic.unwind.ffi.dispose-panic] There are currently no guarantees about the behavior that occurs when a foreign runtime attempts to dispose of, or rethrow, a Rust panic payload. In other words, an unwind originated from a Rust runtime must either lead to termination of the process or be caught by the same runtime. r[link]

# Linkage

[!NOTE] This section is described more in terms of the compiler than of the language.

r[link.intro] The compiler supports various methods to link crates together both statically and dynamically. This section will explore the various methods to link crates together, and more information about native libraries can be found in the <u>FFI section of the book</u>.

r[link.type] In one session of compilation, the compiler can generate multiple artifacts through the usage of either command line flags or the crate\_type attribute. If one or more command line flags are specified, all crate\_type attributes will be ignored in favor of only building the artifacts specified by command line.

r[link.bin]

- --crate-type=bin, #![crate\_type = "bin"] A runnable executable will be produced. This requires that there is a main function in the crate which will be run when the program begins executing. This will link in all Rust and native dependencies, producing a single distributable binary. This is the default crate type.
   r[link.lib]
- --crate-type=lib, #![crate\_type = "lib"] A Rust library will be produced. This is an ambiguous concept as to what exactly is produced because a library can manifest itself in several forms. The purpose of this generic lib option is to generate the "compiler recommended" style of library. The output library will always be usable by rustc, but the actual type of library may change from time-to-time. The remaining output types are all different flavors of libraries, and the lib type can be seen as an alias for one of them (but the actual one is compiler-defined).

r[link.dylib]

 --crate-type=dylib, #![crate\_type = "dylib"] - A dynamic Rust library will be produced. This is different from the lib output type in that this forces dynamic library generation. The resulting dynamic library can be used as a dependency for other libraries and/or executables. This output type will create \*.so files on Linux, \*.dylib files on macOS, and \*.dll files on Windows.

r[link.staticlib]

• --crate-type=staticlib, #![crate\_type = "staticlib"] - A static system library will be produced. This is different from other library outputs in that the compiler will never attempt to link to staticlib outputs. The purpose of this output type is to create a static library containing all of the local crate's code along with all upstream dependencies. This output type will create \*.a files on Linux, macOS and Windows (MinGW), and \*.lib files on Windows (MSVC). This format is recommended for use in situations such as linking Rust code into an existing non-Rust application because it will not have dynamic dependencies on other Rust code.

Note that any dynamic dependencies that the static library may have (such as dependencies on system libraries, or dependencies on Rust libraries that are compiled as dynamic libraries) will have to be specified manually when linking that static library from somewhere. The --print=native-static-libs flag may help with this.

Note that, because the resulting static library contains the code of all the dependencies, including the standard library, and also exports all public symbols of them, linking the static library into an executable or shared library may need special care. In case of a shared library the list of exported symbols will have to be limited via e.g. a linker or symbol version script, exported symbols list (macOS), or module definition file (Windows). Additionally, unused sections can be removed to remove all code of dependencies that is not actually used (e.g. --gc-sections or -dead\_strip for macOS).

r[link.cdylib]

--crate-type=cdylib, #![crate\_type = "cdylib"] - A dynamic system library will be produced. This is used when compiling a dynamic library to be loaded from another language. This output type will create \*.so files on Linux, \*.dylib files on macOS, and \*.dll files on Windows.

r[link.rlib]

 --crate-type=rlib, #![crate\_type = "rlib"] - A "Rust library" file will be produced. This is used as an intermediate artifact and can be thought of as a "static Rust library". These rlib files, unlike staticlib files, are interpreted by the compiler in future linkage. This essentially means that rustc will look for metadata in rlib files like it looks for metadata in dynamic libraries. This form of output is used to produce statically linked executables as well as staticlib outputs.

r[link.proc-macro]

--crate-type=proc-macro, #![crate\_type = "proc-macro"] - The output produced is not specified, but if a -L path is provided to it then the compiler will recognize the output artifacts as a macro and it can be loaded for a program. Crates compiled with this crate type must only export procedural macros. The compiler will automatically set the proc\_macro configuration option. The crates are always compiled with the same target that the compiler itself was built with. For example, if you are executing the compiler from Linux with an x86\_64 CPU, the target will be x86\_64-unknown-linux-gnu even if the crate is a dependency of another crate being built for a different target.

r[link.repetition] Note that these outputs are stackable in the sense that if multiple are specified, then the compiler will produce each form of output without having to recompile. However, this only applies for outputs specified by the same method. If only crate\_type attributes are specified, then they will all be built, but if one or more --crate-type command line flags are specified, then only those outputs will be built.

r[link.dependency] With all these different kinds of outputs, if crate A depends on crate B, then the compiler could find B in various different forms throughout the system. The only forms looked for by the compiler, however, are the rlib format and the dynamic library format. With these two options for a dependent library, the compiler must at some point make a choice between these two formats. With this in mind, the compiler follows these rules when determining what format of dependencies will be used:

r[link.dependency-staticlib]

1. If a static library is being produced, all upstream dependencies are required to be available in rlib formats. This requirement stems from the reason that a dynamic library cannot be converted into a static format.

Note that it is impossible to link in native dynamic dependencies to a static library, and in this case warnings will be printed about all unlinked native dynamic dependencies.

r[link.dependency-rlib]

2. If an rlib file is being produced, then there are no restrictions on what format the upstream dependencies are available in. It is simply required that all upstream dependencies be available for reading metadata from.

The reason for this is that rlib files do not contain any of their upstream dependencies. It wouldn't be very efficient for all rlib files to contain a copy of libstd.rlib!

r[link.dependency-prefer-dynamic]

3. If an executable is being produced and the -C prefer-dynamic flag is not specified, then dependencies are first attempted to be found in the rlib format. If some dependencies are not available in an rlib format, then dynamic linking is attempted (see below).

r[link.dependency-dynamic]

4. If a dynamic library or an executable that is being dynamically linked is being produced, then the compiler will attempt to reconcile the available dependencies in either the rlib or dylib format to create a final product.

A major goal of the compiler is to ensure that a library never appears more than once in any artifact. For example, if dynamic libraries B and C were each statically linked to library A, then a crate could not link to B and C together because there would be two copies of A. The compiler allows mixing the rlib and dylib formats, but this restriction must be satisfied.

The compiler currently implements no method of hinting what format a library should be linked with. When dynamically linking, the compiler will attempt to maximize dynamic dependencies while still allowing some dependencies to be linked in via an rlib.

For most situations, having all libraries available as a dylib is recommended if dynamically linking. For other situations, the compiler will emit a warning if it is unable to determine which formats to link each library with.

In general, --crate-type=bin or --crate-type=lib should be sufficient for all compilation needs, and the other options are just available if more fine-grained control is desired over the output format of a crate.

r[link.crt]

#### **Static and dynamic C runtimes**

r[link.crt.intro] The standard library in general strives to support both statically linked and dynamically linked C runtimes for targets as appropriate. For example the x86\_64-pc-windows-msvc and x86\_64-unknown-linux-musl targets typically come with both runtimes and the user selects which one they'd like. All targets in the compiler have a default mode of linking to the C runtime. Typically targets are linked dynamically by default, but there are exceptions which are static by default such as:

- arm-unknown-linux-musleabi
- arm-unknown-linux-musleabihf
- armv7-unknown-linux-musleabihf
- i686-unknown-linux-musl
- x86\_64-unknown-linux-musl

r[link.crt.crt-static] The linkage of the C runtime is configured to respect the crt-static target feature. These target features are typically configured from the command line via flags to the compiler itself. For example to enable a static runtime you would execute:

```
rustc -C target-feature=+crt-static foo.rs
```

```
whereas to link dynamically to the C runtime you would execute:
rustc -C target-feature=-crt-static foo.rs
```

r[link.crt.ineffective] Targets which do not support switching between linkage of the C runtime will ignore this flag. It's recommended to inspect the resulting binary to ensure that it's linked as you would expect after the compiler succeeds.

r[link.crt.target\_feature] Crates may also learn about how the C runtime is being linked. Code on MSVC, for example, needs to be compiled differently (e.g. with /MT or /MD) depending on the runtime being linked. This is exported currently through the <u>cfg\_attribute\_target\_feature</u> <u>option</u>:

```
#[cfg(target_feature = "crt-static")]
fn foo() {
```

```
println!("the C runtime should be statically linked");
}
#[cfg(not(target_feature = "crt-static"))]
fn foo() {
    println!("the C runtime should be dynamically linked");
}
```

Also note that Cargo build scripts can learn about this feature through <u>environment variables</u>. In a build script you can detect the linkage via:

To use this feature locally, you typically will use the RUSTFLAGS environment variable to specify flags to the compiler through Cargo. For example to compile a statically linked binary on MSVC you would execute: RUSTFLAGS='-C target-feature=+crt-static' cargo build --target x86\_64-pc-windows-msvc

r[link.foreign-code]

### **Mixed Rust and foreign codebases**

r[link.foreign-code.foreign-linkers] If you are mixing Rust with foreign code (e.g. C, C++) and wish to make a single binary containing both types of code, you have two approaches for the final binary link:

- Use rustc. Pass any non-Rust libraries using -L <directory> and -l<library> rustc arguments, and/or #[link] directives in your Rust code. If you need to link against .o files you can use -Clink-arg=file.o.
- Use your foreign linker. In this case, you first need to generate a Rust staticlib target and pass that into your foreign linker invocation. If you need to link multiple Rust subsystems, you will need to generate a *single* staticlib perhaps using lots of extern crate statements to include multiple Rust rlibs. Multiple Rust staticlib files are likely to conflict.

Passing rlibs directly into your foreign linker is currently unsupported.

[!NOTE] Rust code compiled or linked with a different instance of the Rust runtime counts as "foreign code" for the purpose of this section.

r[link.unwinding]

#### **Prohibited linkage and unwinding**

r[link.unwinding.intro] Panic unwinding can only be used if the binary is built consistently according to the following rules.

r[link.unwinding.potential] A Rust artifact is called *potentially unwinding* if any of the following conditions is met:

- The artifact uses the [unwind panic handler][panic.panic\_handler].
- The artifact contains a crate built with the unwind <u>panic strategy</u> that makes a call to a function using a -unwind ABI.
- The artifact makes a "Rust" ABI call to code running in another Rust artifact that has a separate copy of the Rust runtime, and that other

artifact is potentially unwinding.

[!NOTE] This definition captures whether a "Rust" ABI call inside a Rust artifact can ever unwind.

r[link.unwinding.prohibited] If a Rust artifact is potentially unwinding, then all its crates must be built with the unwind <u>panic strategy</u>. Otherwise, unwinding can cause undefined behavior.

[!NOTE] If you are using **rustc** to link, these rules are enforced automatically. If you are *not* using **rustc** to link, you must take care to ensure that unwinding is handled consistently across the entire binary. Linking without **rustc** includes using **dlopen** or similar facilities where linking is done by the system runtime without **rustc** being involved. This can only happen when mixing code with different <u>-C panic</u> flags, so most users do not have to be concerned about this.

[!NOTE] To guarantee that a library will be sound (and linkable with rustc) regardless of the panic runtime used at link-time, the ffi unwind calls lint may be used. The lint flags any calls to - unwind foreign functions or function pointers.

r[asm]

### **Inline assembly**

r[asm.intro] Support for inline assembly is provided via the <u>asm!</u>, <u>naked asm!</u>, and <u>global asm!</u> macros. It can be used to embed handwritten assembly in the assembly output generated by the compiler.

r[asm.stable-targets] Support for inline assembly is stable on the following architectures:

- x86 and x86-64
- ARM
- AArch64 and Arm64EC
- RISC-V
- LoongArch
- s390x

The compiler will emit an error if an assembly macro is used on an unsupported target.

r[asm.example]

#### Example

```
# #[cfg(target_arch = "x86_64")] {
use std::arch::asm;
// Multiply x by 6 using shifts and adds
let mut x: u64 = 4;
unsafe {
    asm!(
        "mov {tmp}, {x}",
        "shl {tmp}, 1",
        "shl {x}, 2",
        "add {x}, {tmp}",
        x = inout(reg) x,
        tmp = out(reg) _,
    );
}
assert_eq!(x, 4 * 6);
# }
```

r[asm.syntax]

#### **Syntax**

```
The following grammar specifies the arguments that can be passed to the
asm!, global_asm! and naked_asm! macros.
```

```
@root AsmArgs -> FormatString (`,` FormatString)* (`,` AsmOperand)*
`,`?
```

```
FormatString -> STRING_LITERAL | RAW_STRING_LITERAL
                                                     MacroInvocation
```

```
AsmOperand ->
```

```
ClobberAbi
```

```
| AsmOptions
```

```
| RegOperand
```

```
ClobberAbi -> `clobber_abi` `(` Abi (`,` Abi)* `,`? `)`
```

```
AsmOptions ->
```

```
`options` `(` ( AsmOption (`,` AsmOption)* `,`? )? `)`
```

```
AsmOption ->
```

```
`pure`
```

```
| `nomem`
```

```
| `readonly`
```

```
> `preserves_flags`
```

```
| `noreturn`
```

```
| `nostack`
```

```
| `att_syntax`
```

)

```
| `raw`
```

```
RegOperand -> (ParamName `=`)?
    (
```

| `label` `{` Statements? `}`

```
DirSpec `(` RegSpec `)` Expression
| DualDirSpec `(` RegSpec `)` DualDirSpecExpression
| `sym` PathExpression
```

```
| `const` Expression
```

```
ParamName -> IDENTIFIER_OR_KEYWORD | RAW_IDENTIFIER
DualDirSpecExpression ->
    Expression
    | Expression `=>` Expression
RegSpec -> RegisterClass | ExplicitRegister
RegisterClass -> IDENTIFIER_OR_KEYWORD
ExplicitRegister -> STRING_LITERAL
DirSpec ->
    `in`
    | `out`
    | `lateout`
DualDirSpec ->
    `inout`
    | `inlateout`
    r[asm.scope]
```
### Scope

r[asm.scope.intro] Inline assembly can be used in one of three ways.

r[asm.scope.asm] With the asm! macro, the assembly code is emitted in a function scope and integrated into the compiler-generated assembly code of a function. This assembly code must obey <u>strict rules</u> to avoid undefined behavior. Note that in some cases the compiler may choose to emit the assembly code as a separate function and generate a call to it.

```
# #[cfg(target_arch = "x86_64")] {
unsafe { core::arch::asm!("/* {} */", in(reg) 0); }
# }
```

r[asm.scope.naked\_asm] With the naked\_asm! macro, the assembly code is emitted in a function scope and constitutes the full assembly code of a function. The naked\_asm! macro is only allowed in <u>naked functions</u>.

```
# #[cfg(target_arch = "x86_64")] {
# #[unsafe(naked)]
# extern "C" fn wrapper() {
core::arch::naked_asm!("/* {} */", const 0);
# }
# }
```

r[asm.scope.global\_asm] With the global\_asm! macro, the assembly code is emitted in a global scope, outside a function. This can be used to hand-write entire functions using assembly code, and generally provides much more freedom to use arbitrary registers and assembler directives.

```
# fn main() {}
# #[cfg(target_arch = "x86_64")]
core::arch::global_asm!("/* {} */", const 0);
```

r[asm.ts-args]

#### **Template string arguments**

r[asm.ts-args.syntax] The assembler template uses the same syntax as <u>format</u> <u>strings</u> (i.e. placeholders are specified by curly braces).

r[asm.ts-args.order] The corresponding arguments are accessed in order, by index, or by name.

```
# #[cfg(target_arch = "x86_64")] {
let x: i64;
let y: i64;
let z: i64;
// This
unsafe { core::arch::asm!("mov {}, {}", out(reg) x, in(reg) 5); }
// ... this
unsafe { core::arch::asm!("mov {0}, {1}", out(reg) y, in(reg) 5); }
// ... and this
unsafe { core::arch::asm!("mov {out}, {in}", out = out(reg) z, in =
in(reg) 5); }
// all have the same behavior
assert_eq!(x, y);
assert_eq!(y, z);
# }
```

r[asm.ts-args.no-implicit] However, implicit named arguments (introduced by <u>RFC #2795</u>) are not supported.

```
# #[cfg(target_arch = "x86_64")] {
let x = 5;
// We can't refer to `x` from the scope directly, we need an operand
like `in(reg) x`
unsafe { core::arch::asm!("/* {x} */"); } // ERROR: no argument
named x
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.ts-args.one-or-more] An asm! invocation may have one or more template string arguments; an asm! with multiple template string arguments is treated as if all the strings were concatenated with a \n between them. The expected usage is for each template string argument to correspond to a line of assembly code.

```
# #[cfg(target_arch = "x86_64")] {
let x: i64;
let y: i64;
// We can separate multiple strings as if they were written
together
unsafe { core::arch::asm!("mov eax, 5", "mov ecx, eax", out("rax")
x, out("rcx") y); }
assert_eq!(x, y);
# }
```

r[asm.ts-args.before-other-args] All template string arguments must appear before any other arguments.

```
let x = 5;
# #[cfg(target_arch = "x86_64")] {
// The template strings need to appear first in the asm invocation
unsafe { core::arch::asm!("/* {x} */", x = const 5, "ud2"); } //
ERROR: unexpected token
# }
# [cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.ts-args.positional-first] As with format strings, positional arguments must appear before named arguments and explicit <u>register operands</u>.

```
# #[cfg(target_arch = "x86_64")] {
```

```
// Named operands need to come after positional ones
unsafe { core::arch::asm!("/* {x} {} */", x = const 5, in(reg) 5); }
```

```
// ERROR: positional arguments cannot follow named arguments or
explicit register arguments
```

```
# }
```

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

```
# #[cfg(target_arch = "x86_64")] {
```

```
// We also can't put explicit registers before positional operands
unsafe { core::arch::asm!("/* {} */", in("eax") 0, in(reg) 5); }
// ERROR: positional arguments cannot follow named arguments or
explicit register arguments
```

# }

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.ts-args.register-operands] Explicit register operands cannot be used by placeholders in the template string.

```
# #[cfg(target_arch = "x86_64")] {
// Explicit register operands don't get substituted, use `eax`
explicitly in the string
unsafe { core::arch::asm!("/* {} */", in("eax") 5); }
// ERROR: invalid reference to argument at index 0
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
r[asm.ts-args.at-least-once] All other named and positional operands must
```

r[asm.ts-args.at-least-once] All other named and positional operands must appear at least once in the template string, otherwise a compiler error is generated. # #[cfg(target\_arch = "x86\_64")] { // We have to name all of the operands in the format string

```
unsafe { core::arch::asm!("", in(reg) 5, x = const 5); }
```

```
// ERROR: multiple unused asm arguments
```

```
# }
```

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.ts-args.opaque] The exact assembly code syntax is target-specific and opaque to the compiler except for the way operands are substituted into the template string to form the code passed to the assembler.

r[asm.ts-args.llvm-syntax] Currently, all supported targets follow the assembly code syntax used by LLVM's internal assembler which usually corresponds to that of the GNU assembler (GAS). On x86, the <code>.intel\_syntax noprefix</code> mode of GAS is used by default. On ARM, the <code>.syntax unified</code> mode is used. These targets impose an additional restriction on the assembly code: any assembler state (e.g. the current section which can be changed with <code>.section</code>) must be restored to its original value at the end of the asm string. Assembly code that does not conform to the GAS syntax will result in assembler-specific behavior. Further constraints on the directives used by inline assembly are indicated by <u>Directives Support</u>.

r[asm.operand-type]

# **Operand type**

r[asm.operand-type.supported-operands] Several types of operands are supported:

r[asm.operand-type.supported-operands.in]

```
in(<reg>) <expr>
```

- <reg> can refer to a register class or an explicit register. The allocated register name is substituted into the asm template string.
- The allocated register will contain the value of <expr> at the start of the assembly code.
- The allocated register must contain the same value at the end of the assembly code (except if a lateout is allocated to the same register).

```
# #[cfg(target_arch = "x86_64")] {
// ``in` can be used to pass values into inline assembly...
unsafe { core::arch::asm!("/* {} */", in(reg) 5); }
# }
```

r[asm.operand-type.supported-operands.out]

```
• out(<reg>) <expr>
```

- <reg> can refer to a register class or an explicit register. The allocated register name is substituted into the asm template string.
- The allocated register will contain an undefined value at the start of the assembly code.
- <expr> must be a (possibly uninitialized) place expression, to which the contents of the allocated register are written at the end of the assembly code.
- An underscore (\_\_) may be specified instead of an expression, which will cause the contents of the register to be discarded at the end of the assembly code (effectively acting as a clobber).

```
# #[cfg(target_arch = "x86_64")] {
let x: i64;
// and `out` can be used to pass values back to rust.
unsafe { core::arch::asm!("/* {} */", out(reg) x); }
# }
```

r[asm.operand-type.supported-operands.lateout]

```
    lateout(<reg>) <expr>
```

- Identical to out except that the register allocator can reuse a register allocated to an in.
- You should only write to the register after all inputs are read, otherwise you may clobber an input.

```
# #[cfg(target_arch = "x86_64")] {
let x: i64;
// `lateout` is the same as `out`
// but the compiler knows we don't care about the value of any
inputs by the
// time we overwrite it.
unsafe { core::arch::asm!("mov {}, 5", lateout(reg) x); }
assert_eq!(x, 5)
# }
```

r[asm.operand-type.supported-operands.inout]

```
    inout(<reg>) <expr>
```

- <reg> can refer to a register class or an explicit register. The allocated register name is substituted into the asm template string.
- The allocated register will contain the value of <expr> at the start of the assembly code.
- <expr> must be a mutable initialized place expression, to which the contents of the allocated register are written at the end of the assembly code.

```
# #[cfg(target_arch = "x86_64")] {
let mut x: i64 = 4;
// `inout` can be used to modify values in-register
unsafe { core::arch::asm!("inc {}", inout(reg) x); }
assert_eq!(x, 5);
# }
```

r[asm.operand-type.supported-operands.inout-arrow]

• inout(<reg>) <in expr> => <out expr>

- Same as inout except that the initial value of the register is taken from the value of <in expr>.
- <out expr> must be a (possibly uninitialized) place expression, to which the contents of the allocated register are written at the end of the assembly code.
- An underscore (\_) may be specified instead of an expression for <out expr>, which will cause the contents of the register to be discarded at the end of the assembly code (effectively acting as a clobber).

• <in expr> and <out expr> may have different types.

```
# #[cfg(target_arch = "x86_64")] {
let x: i64;
// `inout` can also move values to different places
unsafe { core::arch::asm!("inc {}", inout(reg) 4u64=>x); }
assert_eq!(x, 5);
# }
```

r[asm.operand-type.supported-operands.inlateout]

- inlateout(<reg>) <expr> / inlateout(<reg>) <in expr> => <out</li>
  - Identical to inout except that the register allocator can reuse a register allocated to an in (this can happen if the compiler knows the in has the same initial value as the inlateout).
  - You should only write to the register after all inputs are read, otherwise you may clobber an input.

```
# #[cfg(target_arch = "x86_64")] {
let mut x: i64 = 4;
// `inlateout` is `inout` using `lateout`
unsafe { core::arch::asm!("inc {}", inlateout(reg) x); }
assert_eq!(x, 5);
# }
```

r[asm.operand-type.supported-operands.sym]

• sym <path>

• <path> must refer to a fn or static.

- A mangled symbol name referring to the item is substituted into the asm template string.
- The substituted string does not include any modifiers (e.g. GOT, PLT, relocations, etc).
- <path> is allowed to point to a #[thread\_local] static, in which case the assembly code can combine the symbol with relocations (e.g. @plt, @TPOFF) to read from thread-local data.

```
# #[cfg(target_arch = "x86_64")] {
extern "C" fn foo() {
    println!("Hello from inline assembly")
}
// `sym` can be used to refer to a function (even if it doesn't
have an
// external name we can directly write)
unsafe { core::arch::asm!("call {}", sym foo, clobber_abi("C")); }
# }
```

- const <expr>
  - <expr> must be an integer constant expression. This expression follows the same rules as inline const blocks.
  - The type of the expression may be any integer type, but defaults to i32 just like integer literals.
  - The value of the expression is formatted as a string and substituted directly into the asm template string.

```
# #[cfg(target_arch = "x86_64")] {
// swizzle [0, 1, 2, 3] => [3, 2, 0, 1]
const SHUFFLE: u8 = 0b01 00 10 11;
                core::arch::x86_64::__m128 =
let
        x:
                                                      unsafe
                                                                 {
core::mem::transmute([0u32, 1u32, 2u32, 3u32]) };
let y: core::arch::x86_64::__m128;
// Pass a constant value into an instruction that expects
                                                                an
immediate like `pshufd`
unsafe {
    core::arch::asm!("pshufd {xmm}, {xmm}, {shuffle}",
       xmm = inlateout(xmm_reg) x=>y,
        shuffle = const SHUFFLE
```

```
);
}
let y: [u32; 4] = unsafe { core::mem::transmute(y) };
assert_eq!(y, [3, 2, 0, 1]);
# }
```

r[asm.operand-type.supported-operands.label]

```
• label <block>
```

- The address of the block is substituted into the asm template string. The assembly code may jump to the substituted address.
- For targets that distinguish between direct jumps and indirect jumps (e.g. x86-64 with cf-protection enabled), the assembly code must not jump to the substituted address indirectly.
- After execution of the block, the asm! expression returns.
- The type of the block must be unit or ! (never).
- The block starts a new safety context; unsafe operations within the label block must be wrapped in an inner unsafe block, even though the entire asm! expression is already wrapped in unsafe.

```
# #[cfg(target_arch = "x86_64")]
unsafe {
    core::arch::asm!("jmp {}", label {
        println!("Hello from inline assembly label");
    });
}
```

r[asm.operand-type.left-to-right] Operand expressions are evaluated from left to right, just like function call arguments. After the asm! has executed, outputs are written to in left to right order. This is significant if two outputs point to the same place: that place will contain the value of the rightmost output.

```
# #[cfg(target_arch = "x86_64")] {
let mut y: i64;
// y gets its value from the second output, rather than the first
unsafe { core::arch::asm!("mov {}, 0", "mov {}, 1", out(reg) y,
out(reg) y); }
assert_eq!(y, 1);
# }
```

r[asm.operand-type.naked\_asm-restriction] Because naked\_asm! defines a
whole function body and the compiler cannot emit any additional code to handle
operands, it can only use sym and const operands.

r[asm.operand-type.global\_asm-restriction] Because global\_asm! exists outside a function, it can only use sym and const operands.

```
# fn main() {}
// register operands aren't allowed, since we aren't in a function
# #[cfg(target_arch = "x86_64")]
core::arch::global_asm!("", in(reg) 5);
// ERROR: the `in` operand cannot be used with `global_asm!`
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
# fn main() {}
fn foo() {}
# #[cfg(target_arch = "x86_64")]
```

```
// `const` and `sym` are both allowed, however
```

```
core::arch::global_asm!("/* {} {} */", const 0, sym foo);
```

r[asm.register-operands]

### **Register operands**

r[asm.register-operands.register-or-class] Input and output operands can be specified either as an explicit register or as a register class from which the register allocator can select a register. Explicit registers are specified as string literals (e.g. "eax") while register classes are specified as identifiers (e.g. reg).

```
# #[cfg(target_arch = "x86_64")] {
let mut y: i64;
// We can name both `reg`, or an explicit register like `eax` to
get an
// integer register
unsafe { core::arch::asm!("mov eax, {:e}", in(reg) 5,
lateout("eax") y); }
assert_eq!(y, 5);
# }
```

r[asm.register-operands.equivalence-to-base-register] Note that explicit registers treat register aliases (e.g. r14 vs lr on ARM) and smaller views of a register (e.g. eax vs rax) as equivalent to the base register.

r[asm.register-operands.error-two-operands] It is a compile-time error to use the same explicit register for two input operands or two output operands.

```
# #[cfg(target_arch = "x86_64")] {
// We can't name eax twice
unsafe { core::arch::asm!("", in("eax") 5, in("eax") 4); }
// ERROR: register `eax` conflicts with register `eax`
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
# #[cfg(target_arch = "x86_64")] {
// ... even using different aliases
unsafe { core::arch::asm!("", in("ax") 5, in("rax") 4); }
// ERROR: register `rax` conflicts with register `ax`
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-operands.error-overlapping] Additionally, it is also a compiletime error to use overlapping registers (e.g. ARM VFP) in input operands or in output operands.

```
# #[cfg(target_arch = "x86_64")] {
// al overlaps with ax, so we can't name both of them.
unsafe { core::arch::asm!("", in("ax") 5, in("al") 4i8); }
// ERROR: register `al` conflicts with register `ax`
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-operands.allowed-types] Only the following types are allowed as operands for inline assembly:

- Integers (signed and unsigned)
- Floating-point numbers
- Pointers (thin only)
- Function pointers
- SIMD vectors (structs defined with #[repr(simd)] and which implement Copy). This includes architecture-specific vector types defined in
   atducereb such as m120 (x26) or integrate t (ADM)

```
std::arch such as __m128 (x86) or int8x16_t (ARM).
```

```
# #[cfg(target_arch = "x86_64")] {
extern "C" fn foo() {}
// Integers are allowed...
let y: i64 = 5;
unsafe { core::arch::asm!("/* {} */", in(reg) y); }
// and pointers...
let py = &raw const y;
unsafe { core::arch::asm!("/* {} */", in(reg) py); }
// floats as well...
let f = 1.0f32;
unsafe { core::arch::asm!("/* {} */", in(xmm_reg) f); }
// even function pointers and simd vectors.
let func: extern "C" fn() = foo;
unsafe { core::arch::asm!("/* {} */", in(reg) func); }
let z = unsafe { core::arch::x86_64::_mm_set_epi64x(1, 0) };
```

```
unsafe { core::arch::asm!("/* {} */", in(xmm_reg) z); }
# }
# #[cfg(target_arch = "x86_64")] {
struct Foo;
let x: Foo = Foo;
// Complex types like structs are not allowed
unsafe { core::arch::asm!("/* {} */", in(reg) x); }
// ERROR: cannot use value of type `Foo` for inline assembly
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-operands.supported-register-classes] Here is the list of currently supported register classes:

Architecture	Register class	Registers	LLVM constraint code
x86	reg	ax, bx, cx, dx, si, di, bp, r[8-15] (x86- 64 only)	r
x86	reg_ab cd	ax, bx, cx, dx	Q
x86-32	reg_by te	al, bl, cl, dl, ah, bh, ch, dh	q
x86-64	reg_by te*	al, bl, cl, dl, sil, dil, bpl, r[8-15]b	q
x86	xmm_re g	xmm[0-7] (x86) xmm[0- 15] (x86-64)	X
x86	ymm_re g	ymm[0-7] (x86) ymm[0- 15] (x86-64)	X
x86	zmm_re g	zmm[0-7] (x86) zmm[0- 31] (x86-64)	V
x86	kreg	k[1-7]	Yk

Architecture	Register class	Registers	LLVM constraint code
x86	kreg0	k0	Only clobbers
x86	x87_re g	st([0-7])	Only clobbers
x86	mmx_re g	mm[0-7]	Only clobbers
x86-64	tmm_re g	tmm[0-7]	Only clobbers
AArch64	reg	×[0-30]	r
AArch64	vreg	v[0-31]	W
AArch64	vreg_l ow16	v[0-15]	X
AArch64	preg	p[0-15], ffr	Only clobbers
Arm64EC	reg	x[0-12], x[15-22], x[25-27], x30	r
Arm64EC	vreg	v[0-15]	W
Arm64EC	vreg_l ow16	v[0-15]	X
ARM (ARM/Thumb2)	reg	r[0-12], r14	r
ARM (Thumb1)	reg	r[0-7]	r
ARM	sreg	s[0-31]	t
ARM	sreg_l ow16	s[0-15]	X
ARM	dreg	d[0-31]	W

Architecture	Register class	Registers	LLVM constraint code
ARM	dreg_l ow16	d[0-15]	t
ARM	dreg_l ow8	d[0-8]	X
ARM	qreg	q[0-15]	W
ARM	qreg_l ow8	q[0-7]	t
ARM	qreg_l ow4	q[0-3]	X
RISC-V	reg	x1, x[5-7], x[9-15], x[16-31] (non-RV32E)	r
RISC-V	freg	f[0-31]	f
RISC-V	vreg	v[0-31]	Only clobbers
LoongArch	reg	<pre>\$r1, \$r[4-20], \$r[23,30]</pre>	r
LoongArch	freg	\$f[0-31]	f
s390x	reg	r[0-10], r[12-14]	r
s390x	reg_ad dr	r[1-10], r[12-14]	a
s390x	freg	f[0-15]	f
s390x	vreg	v[0-31]	Only clobbers
s390x	areg	a[2-15]	Only clobbers

[!NOTE]

- On x86 we treat reg\_byte differently from reg because the compiler can allocate al and an separately whereas reg reserves the whole register.
- On x86-64 the high byte registers (e.g. ah) are not available in the reg\_byte register class.
- Some register classes are marked as "Only clobbers" which means that registers in these classes cannot be used for inputs or outputs, only clobbers of the form out(<explicit register>) \_ or lateout(<explicit register>) \_.

r[asm.register-operands.value-type-constraints] Each register class has constraints on which value types they can be used with. This is necessary because the way a value is loaded into a register depends on its type. For example, on bigendian systems, loading a i32×4 and a i8×16 into a SIMD register may result in different register contents even if the byte-wise memory representation of both values is identical. The availability of supported types for a particular register class may depend on what target features are currently enabled.

Architecture	Register class	Target feature	Allowed types
x86-32	reg	None	i16, i32, f32
x86-64	reg	None	i16, i32, f32, i64, f64
x86	reg_by te	None	18
x86	xmm_re g	sse	i32, f32, i64, f64, i8x16, i16x8, i32x4, i64x2, f32x4, f64x2
x86	ymm_re g	avx	i32, f32, i64, f64, i8x16, i16x8, i32x4, i64x2, f32x4, f64x2 i8x32, i16x16, i32x8, i64x4, f32x8, f64x4

Architecture	Register class	Target feature	Allowed types
x86	zmm_re g	avx51 2f	<pre>i32, f32, i64, f64, i8x16, i16x8, i32x4, i64x2, f32x4, f64x2 i8x32, i16x16, i32x8, i64x4, f32x8, f64x4 i8x64, i16x32, i32x16, i64x8, f32x16, f64x8</pre>
x86	kreg	avx51 2f	i8, i16
x86	kreg	avx51 2bw	i32, i64
x86	mmx_re g	N/A	Only clobbers
x86	x87_re g	N/A	Only clobbers
x86	tmm_re g	N/A	Only clobbers
AArch64	reg	None	i8, i16, i32, f32, i64, f64
AArch64	vreg	neon	<pre>i8, i16, i32, f32, i64, f64, i8x8, i16x4, i32x2, i64x1, f32x2, f64x1, i8x16, i16x8, i32x4, i64x2, f32x4, f64x2</pre>
AArch64	preg	N/A	Only clobbers
Arm64EC	reg	None	i8, i16, i32, f32, i64, f64

Architecture	Register class	Target feature	Allowed types
Arm64EC	vreg	neon	<pre>i8, i16, i32, f32, i64, f64, i8x8, i16x4, i32x2, i64x1, f32x2, f64x1, i8x16, i16x8, i32x4, i64x2, f32x4, f64x2</pre>
ARM	reg	None	i8, i16, i32, f32
ARM	sreg	vfp2	i32, f32
ARM	dreg	vfp2	i64, f64, i8x8, i16x4, i32x2, i64x1, f32x2
ARM	qreg	neon	i8x16, i16x8, i32x4, i64x2,f32x4
RISC-V32	reg	None	i8, i16, i32, f32
RISC-V64	reg	None	i8, i16, i32, f32, i64, f64
RISC-V	freg	f	f32
RISC-V	freg	d	f64
RISC-V	vreg	N/A	Only clobbers
LoongArch64	reg	None	i8, i16, i32, i64, f32, f64
LoongArch64	freg	f	f32
LoongArch64	freg	d	f64
s390x	reg, reg_ad dr	None	i8, i16, i32, i64
s390x	freg	None	f32, f64
s390x	vreg	N/A	Only clobbers
s390x	areg	N/A	Only clobbers

[!NOTE] For the purposes of the above table pointers, function pointers and isize/usize are treated as the equivalent integer type (i16/i32/i64 depending on the target).

```
# #[cfg(target_arch = "x86_64")] {
 let x = 5i32;
 let y = -1i8;
 let z = unsafe { core::arch::x86_64::_mm_set_epi64x(1, 0) };
 // reg is valid for `i32`, `reg_byte` is valid for `i8`, and
 xmm_reg is valid for `__m128i`
 // We can't use `tmmO` as an input or output, but we can clobber
 it.
 unsafe { core::arch::asm!("/* {} {} {} */", in(reg) x, in(reg_byte)
 y, in(xmm_reg) z, out("tmm0") _); }
 # }
# #[cfg(target_arch = "x86_64")] {
let z = unsafe { core::arch::x86_64::_mm_set_epi64x(1, 0) };
// We can't pass an `__m128i` to a `reg` input
unsafe { core::arch::asm!("/* {} */", in(reg) z); }
// ERROR: type `__m128i` cannot be used with this register class
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-operands.smaller-value] If a value is of a smaller size than the register it is allocated in then the upper bits of that register will have an undefined value for inputs and will be ignored for outputs. The only exception is the freg register class on RISC-V where f32 values are NaN-boxed in a f64 as required by the RISC-V architecture.

```
# #[cfg(target_arch = "x86_64")] {
let mut x: i64;
// Moving a 32-bit value into a 64-bit value, oops.
#[allow(asm_sub_register)] // rustc warns about this behavior
unsafe { core::arch::asm!("mov {}, {}", lateout(reg) x, in(reg)
4i32); }
// top 32-bits are indeterminate
assert_eq!(x, 4); // This assertion is not guaranteed to succeed
```

assert\_eq!(x & 0xFFFFFFF, 4); // However, this one will succeed
# }

r[asm.register-operands.separate-input-output] When separate input and output expressions are specified for an inout operand, both expressions must have the same type. The only exception is if both operands are pointers or integers, in which case they are only required to have the same size. This restriction exists because the register allocators in LLVM and GCC sometimes cannot handle tied operands with different types.

```
# #[cfg(target_arch = "x86_64")] {
 // Pointers and integers can mix (as long as they are the same
 size)
 let x: isize = 0;
 let y: *mut ();
 // Transmute an `isize` to a `*mut ()`, using inline assembly magic
 unsafe { core::arch::asm!("/*{}*/", inout(reg) x=>y); }
 assert!(y.is_null()); // Extremely roundabout way to make a null
 pointer
 # }
# #[cfg(target_arch = "x86_64")] {
let x: i32 = 0;
let y: f32;
// But we can't reinterpret an `i32` to an `f32` like this
unsafe { core::arch::asm!("/* {} */", inout(reg) x=>y); }
// ERROR: incompatible types for asm inout argument
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-names]

## **Register names**

r[asm.register-names.supported-register-aliases] Some registers have multiple names. These are all treated by the compiler as identical to the base register name. Here is the list of all supported register aliases:

Architecture	Base register	Aliases
x86	ax	eax, rax
x86	bx	ebx, rbx
x86	CX	ecx, rcx
x86	dx	edx, rdx
x86	si	esi, rsi
x86	di	edi, rdi
x86	bp	bpl, ebp, rbp
x86	sp	spl, esp, rsp
x86	ip	eip, rip
x86	st(0)	st
x86	r[8-15]	r[8-15]b, r[8-15]w, r[8-15]d
x86	xmm[0- 31]	ymm[0-31], zmm[0-31]
AArch64	x[0-30]	w[0-30]
AArch64	x29	fp
AArch64	×30	lr
AArch64	sp	wsp
AArch64	xzr	wzr
AArch64	v[0-31]	b[0-31], h[0-31], s[0-31], d[0-31], q[0-31]
Arm64EC	×[0-30]	w[0-30]

Architecture	Base register	Aliases
Arm64EC	x29	fp
Arm64EC	×30	lr
Arm64EC	sp	wsp
Arm64EC	xzr	WZr
Arm64EC	v[0-15]	b[0-15], h[0-15], s[0-15], d[0-15], q[0-15]
ARM	r[0-3]	a[1-4]
ARM	r[4-9]	v[1-6]
ARM	r9	rfp
ARM	r10	sl
ARM	r11	fp
ARM	r12	ір
ARM	r13	sp
ARM	r14	lr
ARM	r15	pc
RISC-V	×0	zero
RISC-V	X1	ra
RISC-V	x2	sp
RISC-V	<b>x3</b>	gp
RISC-V	x4	tp
RISC-V	x[5-7]	t[0-2]
RISC-V	x8	fp, s0
RISC-V	×9	S1
RISC-V	x[10- 17]	a[0-7]

Architecture	Base register	Aliases
RISC-V	x[18- 27]	s[2-11]
RISC-V	x[28- 31]	t[3-6]
RISC-V	f[0-7]	ft[0-7]
RISC-V	f[8-9]	fs[0-1]
RISC-V	f[10- 17]	fa[0-7]
RISC-V	f[18- 27]	fs[2-11]
RISC-V	f[28- 31]	ft[8-11]
LoongArch	\$r0	\$zero
LoongArch	\$r1	\$ra
LoongArch	\$r2	\$tp
LoongArch	\$r3	\$sp
LoongArch	\$r[4- 11]	\$a[0-7]
LoongArch	\$r[12- 20]	\$t[0-8]
LoongArch	\$r21	
LoongArch	\$r22	\$fp, \$s9
LoongArch	\$r[23- 31]	\$s[0-8]
LoongArch	\$f[0-7]	\$fa[0-7]
LoongArch	\$f[8- 23]	\$ft[0-15]

Architecture	Base register	Aliases
LoongArch	\$f[24-	\$fs[0-7]
	31]	

```
# #[cfg(target_arch = "x86_64")] {
let z = 0i64;
// rax is an alias for eax and ax
unsafe { core::arch::asm!("", in("rax") z); }
# }
```

r[asm.register-names.not-for-io] Some registers cannot be used for input or output operands:

Architecture	Unsupported register	Reason
All	sp, r15 (s390x)	The stack pointer must be restored to its original value at the end of the assembly code or before jumping to a label block.
All	bp(x86),x29(AArch64andArm64EC),x8(RISC-V),\$fp(LoongArch),r11(s390x)	The frame pointer cannot be used as an input or output.
ARM	r7 or r11	On ARM the frame pointer can be either r7 or r11 depending on the target. The frame pointer cannot be used as an input or output.

Architecture	Unsupported register	Reason
All	si (x86-32), bx (x86-64), r6 (ARM), x19 (AArch64 and Arm64EC), x9 (RISC-V), \$s8 (LoongArch)	This is used internally by LLVM as a "base pointer" for functions with complex stack frames.
x86	ip	This is the program counter, not a real register.
AArch64	xzr	This is a constant zero register which can't be modified.
AArch64	x18	This is an OS-reserved register on some AArch64 targets.
Arm64EC	xzr	This is a constant zero register which can't be modified.
Arm64EC	x18	This is an OS-reserved register.
Arm64EC	x13, x14, x23, x24, x28, v[16- 31],p[0-15],ffr	These are AArch64 registers that are not supported for Arm64EC.
ARM	pc	This is the program counter, not a real register.
ARM	r9	This is an OS-reserved register on some ARM targets.
RISC-V	×0	This is a constant zero register which can't be modified.
RISC-V	gp, tp	These registers are reserved and cannot be used as inputs or outputs.

Architecture	Unsupported register	Reason
LoongArch	\$r0 or \$zero	This is a constant zero register which can't be modified.
LoongArch	\$r2 or \$tp	This is reserved for TLS.
LoongArch	\$r21	This is reserved by the ABI.
s390x	c[0-15]	Reserved by the kernel.
s390x	a[0-1]	Reserved for system use.

```
# #[cfg(target_arch = "x86_64")] {
```

// bp is reserved

```
unsafe { core::arch::asm!("", in("bp") 5i32); }
```

```
// ERROR: invalid register `bp`: the frame pointer cannot be used as
an operand for inline asm
# }
```

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.register-names.fp-bp-reserved] The frame pointer and base pointer registers are reserved for internal use by LLVM. While asm! statements cannot explicitly specify the use of reserved registers, in some cases LLVM will allocate one of these reserved registers for reg operands. Assembly code making use of reserved registers should be careful since reg operands may use the same registers.

r[asm.template-modifiers]

## **Template modifiers**

r[asm.template-modifiers.intro] The placeholders can be augmented by modifiers which are specified after the : in the curly braces. These modifiers do not affect register allocation, but change the way operands are formatted when inserted into the template string.

r[asm.template-modifiers.only-one] Only one modifier is allowed per template placeholder.

```
# #[cfg(target_arch = "x86_64")] {
// We can't specify both `r` and `e` at the same time.
unsafe { core::arch::asm!("/* {:er}", in(reg) 5i32); }
// ERROR: asm template modifier must be a single character
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.template-modifiers.supported-modifiers] The supported modifiers are a subset of LLVM's (and GCC's) <u>asm template argument modifiers</u>, but do not use the same letter codes.

Architecture	Register class	Modifier	Example output	LLVM modifier
x86-32	reg	None	eax	k
x86-64	reg	None	rax	q
x86-32	reg_ab cd	l	al	b
x86-64	reg	l	al	b
x86	reg_ab cd	h	ah	h
x86	reg	X	ax	W
x86	reg	е	eax	k
x86-64	reg	r	rax	q
x86	reg_by te	None	al/ah	None

Architecture	Register class	Modifier	Example output	LLVM modifier
x86	xmm_re g	None	×mm0	x
x86	ymm_re g	None	ymm0	t
x86	zmm_re g	None	zmm0	g
x86	*mm_re g	X	×mm0	X
x86	*mm_re g	У	ymm0	t
x86	*mm_re g	Z	zmm0	g
x86	kreg	None	k1	None
AArch64/Arm64EC	reg	None	×0	X
AArch64/Arm64EC	reg	W	w0	W
AArch64/Arm64EC	reg	X	×0	X
AArch64/Arm64EC	vreg	None	V0	None
AArch64/Arm64EC	vreg	V	v0	None
AArch64/Arm64EC	vreg	b	b0	b
AArch64/Arm64EC	vreg	h	h0	h
AArch64/Arm64EC	vreg	S	s0	S
AArch64/Arm64EC	vreg	d	d0	d
AArch64/Arm64EC	vreg	q	q0	q
ARM	reg	None	r0	None
ARM	sreg	None	s0	None
ARM	dreg	None	d0	Ρ
ARM	qreg	None	q0	q

Architecture	Register class	Modifier	Example output	LLVM modifier
ARM	qreg	e / f	d0 / d1	e / f
RISC-V	reg	None	X1	None
RISC-V	freg	None	f0	None
LoongArch	reg	None	\$r1	None
LoongArch	freg	None	\$f0	None
s390x	reg	None	%r0	None
s390x	reg_ad dr	None	%r1	None
s390x	freg	None	%f0	None

[!NOTE]

- on ARM e / f: this prints the low or high doubleword register name of a NEON quad (128-bit) register.
- on x86: our behavior for reg with no modifiers differs from what GCC does. GCC will infer the modifier based on the operand value type, while we default to the full register size.
- on x86 xmm\_reg: the x, t and g LLVM modifiers are not yet implemented in LLVM (they are supported by GCC only), but this should be a simple change.

```
# #[cfg(target_arch = "x86_64")] {
let mut x = 0x10u16;
// u16::swap_bytes using `xchg`
// low half of `{x}` is referred to by `{x:l}`, and the high half
by `{x:h}`
unsafe { core::arch::asm!("xchg {x:l}, {x:h}", x = inout(reg_abcd)
x); }
assert_eq!(x, 0x1000u16);
# }
```

r[asm.template-modifiers.smaller-value] As stated in the previous section, passing an input value smaller than the register width will result in the upper bits

of the register containing undefined values. This is not a problem if the inline asm only accesses the lower bits of the register, which can be done by using a template modifier to use a subregister name in the assembly code (e.g. ax instead of rax). Since this an easy pitfall, the compiler will suggest a template modifier to use where appropriate given the input type. If all references to an operand already have modifiers then the warning is suppressed for that operand.

r[asm.abi-clobbers]

#### **ABI clobbers**

r[asm.abi-clobbers.intro] The clobber\_abi keyword can be used to apply a default set of clobbers to the assembly code. This will automatically insert the necessary clobber constraints as needed for calling a function with a particular calling convention: if the calling convention does not fully preserve the value of a register across a call then lateout("...") \_ is implicitly added to the operands list (where the ... is replaced by the register's name).

```
# #[cfg(target_arch = "x86_64")] {
extern "C" fn foo() -> i32 { 0 }
let z: i32;
// To call a function, we have to inform the compiler that we're
clobbering
// callee saved registers
unsafe { core::arch::asm!("call {}", sym foo, out("rax") z,
clobber_abi("C")); }
assert_eq!(z, 0);
# }
```

r[asm.abi-clobbers.many] clobber\_abi may be specified any number of times. It will insert a clobber for all unique registers in the union of all specified calling conventions.

```
# #[cfg(target_arch = "x86_64")] {
extern "sysv64" fn foo() -> i32 { 0 }
extern "win64" fn bar(x: i32) -> i32 { x + 1}

let z: i32;
// We can even call multiple functions with different conventions
and
// different saved registers
unsafe {
    core::arch::asm!(
        "call {}",
        "mov ecx, eax",
        "call {}",
        sym foo,
        sym bar,
```

```
out("rax") z,
clobber_abi("C")
);
}
assert_eq!(z, 1);
# }
```

r[asm.abi-clobbers.must-specify] Generic register class outputs are disallowed by the compiler when clobber\_abi is used: all outputs must specify an explicit register.

```
# #[cfg(target_arch = "x86_64")] {
extern "C" fn foo(x: i32) -> i32 { 0 }
let z: i32;
// explicit registers must be used to not accidentally overlap.
unsafe {
    core::arch::asm!(
        "mov eax, {:e}",
        "call {}",
        out(reg) z,
        sym foo,
        clobber_abi("C")
    );
    // ERROR: asm with `clobber_abi` must specify explicit registers
for outputs
}
assert_eq!(z, 0);
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.abi-clobbers.explicit-have-precedence] Explicit register outputs have precedence over the implicit clobbers inserted by clobber\_abi: a clobber will only be inserted for a register if that register is not used as an output.

r[asm.abi-clobbers.supported-abis] The following ABIs can be used with
clobber\_abi:

Architecture	ABI name	<b>Clobbered registers</b>
--------------	----------	----------------------------

Architecture	ABI name	<b>Clobbered registers</b>
x86-32	<pre>"C", "system", "efiapi", "cdecl", "stdcall", "fastcall"</pre>	<pre>ax, cx, dx, xmm[0-7], mm[0-7], k[0-7], st([0- 7])</pre>
x86-64	"C", "system" (on Windows), "efiapi", "win64"	ax, cx, dx, r[8-11], xmm[0-31], mm[0-7], k[0- 7],st([0-7]),tmm[0-7]
x86-64	"C", "system" (on non-Windows), "sysv64"	<pre>ax, cx, dx, si, di, r[8- 11], xmm[0-31], mm[0-7], k[0-7], st([0-7]), tmm[0- 7]</pre>
AArch64	"C", "system", "efiapi"	x[0-17], x18*, x30, v[0- 31],p[0-15],ffr
Arm64EC	"C", "system"	x[0-12], x[15-17], x30, v[0-15]
ARM	"C", "system", "efiapi", "aapcs"	r[0-3], r12, r14, s[0- 15],d[0-7],d[16-31]
RISC-V	"C", "system", "efiapi"	x1, x[5-7], x[10-17]*, x[28-31]*, f[0-7], f[10- 17], f[28-31], v[0-31]
LoongArch	"C", "system"	<pre>\$r1, \$r[4-20], \$f[0-23]</pre>
s390x	"C", "system"	r[0-5], r14, f[0-7], v[0- 31], a[2-15]

#### [!NOTE]

- On AArch64 ×18 only included in the clobber list if it is not considered as a reserved register on the target.
- On RISC-V ×[16-17] and ×[28-31] only included in the clobber list if they are not considered as reserved registers on the target.

The list of clobbered registers for each ABI is updated in rustc as architectures gain new registers: this ensures that asm! clobbers will continue to be correct when LLVM starts using these new registers in its generated code.

r[asm.options]

## **Options**

r[asm.options.supported-options] Flags are used to further influence the behavior of the inline assembly code. Currently the following options are defined: r[asm.options.supported-options.pure]

• pure: The assembly code has no side effects, must eventually return, and its outputs depend only on its direct inputs (i.e. the values themselves, not what they point to) or values read from memory (unless the nomem options is also set). This allows the compiler to execute the assembly code fewer times than specified in the program (e.g. by hoisting it out of a loop) or even eliminate it entirely if the outputs are not used. The pure option must be combined with either the nomem or readonly options, otherwise a compile-time error is emitted.

```
# #[cfg(target_arch = "x86_64")] {
 let x: i32 = 0;
 let z: i32;
 // pure can be used to optimize by assuming the assembly has no
 side effects
 unsafe
          {
           core::arch::asm!("inc
                                     {}",
                                            inout(reg)
                                                         Х
                                                            =>
                                                                 z,
 options(pure, nomem)); }
 assert_eq!(z, 1);
 # }
# #[cfg(target_arch = "x86_64")] {
let x: i32 = 0;
let z: i32;
// Either nomem or readonly must be satisfied, to indicate whether
or not
// memory is allowed to be read
             core::arch::asm!("inc {}",
unsafe
         {
                                            inout(req)
                                                        Х
                                                             =>
                                                                  z,
options(pure)); }
// ERROR: the `pure` option must be combined with either `nomem` or
`readonly`
assert_eq!(z, 0);
# }
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.options.supported-options.nomem]

 nomem: The assembly code does not read from or write to any memory accessible outside of the assembly code. This allows the compiler to cache the values of modified global variables in registers across execution of the assembly code since it knows that they are not read from or written to by it. The compiler also assumes that the assembly code does not perform any kind of synchronization with other threads, e.g. via fences.

```
# #[cfg(target_arch = "x86_64")] {
let mut x = 0i32;
let z: i32;
// Accessing outside memory from assembly when `nomem` is
// specified is disallowed
unsafe {
    core::arch::asm!("mov {val:e}, dword ptr [{ptr}]",
        ptr = in(reg) \& mut x,
        val = lateout(reg) z,
        options(nomem)
    )
}
// Writing to outside memory from assembly when `nomem` is
// specified is also undefined behaviour
unsafe {
    core::arch::asm!("mov dword ptr [{ptr}], {val:e}",
        ptr = in(reg) \& mut x,
        val = in(reg) z,
        options(nomem)
    )
}
# }
 # #[cfg(target_arch = "x86_64")] {
 let x: i32 = 0;
 let z: i32;
 // If we allocate our own memory, such as via `push`, however.
 // we can still use it
 unsafe {
      core::arch::asm!("push {x}", "add qword ptr [rsp], 1", "pop
```
r[asm.options.supported-options.readonly]

• readonly: The assembly code does not write to any memory accessible outside of the assembly code. This allows the compiler to cache the values of unmodified global variables in registers across execution of the assembly code since it knows that they are not written to by it. The compiler also assumes that this assembly code does not perform any kind of synchronization with other threads, e.g. via fences.

```
# #[cfg(target_arch = "x86_64")] {
let mut x = 0;
// We cannot modify outside memory when `readonly` is specified
unsafe {
        core::arch::asm!("mov dword ptr[{}], 1", in(reg) &mut x,
options(readonly))
}
# }
 # #[cfg(target_arch = "x86_64")] {
 let x: i64 = 0;
 let z: i64;
 // We can still read from it, though
 unsafe {
     core::arch::asm!("mov {x}, qword ptr [{x}]",
         x = inout(reg) \& x \Rightarrow z,
         options(readonly)
     );
 }
 assert_eq!(z, 0);
 # }
 # #[cfg(target_arch = "x86_64")] {
 let x: i64 = 0;
 let z: i64;
```

```
// Same exception applies as with nomem.
unsafe {
    core::arch::asm!("push {x}", "add qword ptr [rsp], 1", "pop
{x}",
    x = inout(reg) x => z,
    options(readonly)
    );
}
assert_eq!(z, 1);
# }
```

r[asm.options.supported-options.preserves\_flags]

• preserves\_flags: The assembly code does not modify the flags register (defined in the rules below). This allows the compiler to avoid recomputing the condition flags after execution of the assembly code.

r[asm.options.supported-options.noreturn]

noreturn: The assembly code does not fall through; behavior is undefined if it does. It may still jump to label blocks. If any label blocks return unit, the asm! block will return unit. Otherwise it will return ! (never). As with a call to a function that does not return, local variables in scope are not dropped before execution of the assembly code.

```
fn main() -> ! {
# #[cfg(target_arch = "x86_64")] {
     // We can use an instruction to trap execution inside of a
noreturn block
    unsafe { core::arch::asm!("ud2", options(noreturn)); }
# }
# #[cfg(not(target_arch = "x86_64"))] panic!("no return");
}
# #[cfg(target_arch = "x86_64")] {
// You are responsible for not falling past the end of a noreturn
asm block
unsafe { core::arch::asm!("", options(noreturn)); }
# }
 # #[cfg(target_arch = "x86_64")]
 let _: () = unsafe {
     // You may still jump to a `label` block
```

```
core::arch::asm!("jmp {}", label {
    println!();
  }, options(noreturn));
};
```

r[asm.options.supported-options.nostack]

• **nostack**: The assembly code does not push data to the stack, or write to the stack red-zone (if supported by the target). If this option is *not* used then the stack pointer is guaranteed to be suitably aligned (according to the target ABI) for a function call.

```
# #[cfg(target_arch = "x86_64")] {
// `push` and `pop` are UB when used with nostack
unsafe { core::arch::asm!("push rax", "pop rax", options(nostack));
}
# }
```

r[asm.options.supported-options.att\_syntax]

 att\_syntax: This option is only valid on x86, and causes the assembler to use the .att\_syntax prefix mode of the GNU assembler. Register operands are substituted in with a leading %.

```
# #[cfg(target_arch = "x86_64")] {
let x: i32;
let y = 1i32;
// We need to use AT&T Syntax here. src, dest order for operands
unsafe {
    core::arch::asm!("mov {y:e}, {x:e}",
        x = lateout(reg) x,
        y = in(reg) y,
        options(att_syntax)
    );
}
assert_eq!(x, y);
# }
```

r[asm.options.supported-options.raw]

raw: This causes the template string to be parsed as a raw assembly string, with no special handling for { and }. This is primarily useful when

including raw assembly code from an external file using include\_str!.

r[asm.options.checks] The compiler performs some additional checks on options:

r[asm.options.checks.mutually-exclusive]

```
• The nomem and readonly options are mutually exclusive: it is a compile-
time error to specify both.
```

```
# #[cfg(target_arch = "x86_64")] {
```

// nomem is strictly stronger than readonly, they can't be specified
together

```
unsafe { core::arch::asm!("", options(nomem, readonly)); }
```

```
// ERROR: the `nomem` and `readonly` options are mutually exclusive
# }
```

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.options.checks.pure]

• It is a compile-time error to specify pure on an asm block with no outputs or only discarded outputs (\_).

```
# #[cfg(target_arch = "x86_64")] {
```

```
// pure blocks need at least one output
```

```
unsafe { core::arch::asm!("", options(pure)); }
```

```
// ERROR: asm with the `pure` option must have at least one output
# }
```

```
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

r[asm.options.checks.noreturn]

```
    It is a compile-time error to specify noreturn on an asm block with outputs
and without labels.
    # #[cfg(target_arch = "x86_64")] {
```

```
let z: i32;
// noreturn can't have outputs
unsafe { core::arch::asm!("mov {:e}, 1", out(reg) z,
options(noreturn)); }
// ERROR: asm outputs are not allowed with the `noreturn` option
# }
```

# #[cfg(not(target\_arch = "x86\_64"))] core::compile\_error!("Test not supported on this arch");

r[asm.options.checks.label-with-outputs]

• It is a compile-time error to have any label blocks in an asm block with outputs.

r[asm.options.naked\_asm-restriction] naked\_asm! only supports the att\_syntax and raw options. The remaining options are not meaningful because the inline assembly defines the whole function body.

r[asm.options.global\_asm-restriction] global\_asm! only supports the
att\_syntax and raw options. The remaining options are not meaningful for
global-scope inline assembly.

```
# fn main() {}
# #[cfg(target_arch = "x86_64")]
// nomem is useless on global_asm!
core::arch::global_asm!("", options(nomem));
# #[cfg(not(target_arch = "x86_64"))] core::compile_error!("Test not
supported on this arch");
```

```
r[asm.rules]
```

## **Rules for inline assembly**

r[asm.rules.intro] To avoid undefined behavior, these rules must be followed when using function-scope inline assembly (asm!):

r[asm.rules.reg-not-input]

- Any registers not specified as inputs will contain an undefined value on entry to the assembly code.
  - An "undefined value" in the context of inline assembly means that the register can (non-deterministically) have any one of the possible values allowed by the architecture. Notably it is not the same as an LLVM undef which can have a different value every time you read it (since such a concept does not exist in assembly code).

r[asm.rules.reg-not-output]

- Any registers not specified as outputs must have the same value upon exiting the assembly code as they had on entry, otherwise behavior is undefined.
  - This only applies to registers which can be specified as an input or output. Other registers follow target-specific rules.
  - Note that a lateout may be allocated to the same register as an in, in which case this rule does not apply. Code should not rely on this however since it depends on the results of register allocation.

r[asm.rules.unwind]

- Behavior is undefined if execution unwinds out of the assembly code.
  - This also applies if the assembly code calls a function which then unwinds.

r[asm.rules.mem-same-as-ffi]

- The set of memory locations that assembly code is allowed to read and write are the same as those allowed for an FFI function.
  - If the readonly option is set, then only memory reads are allowed.
  - If the nomem option is set then no reads or writes to memory are allowed.

• These rules do not apply to memory which is private to the assembly code, such as stack space allocated within it.

r[asm.rules.black-box]

- The compiler cannot assume that the instructions in the assembly code are the ones that will actually end up executed.
  - This effectively means that the compiler must treat the assembly code as a black box and only take the interface specification into account, not the instructions themselves.
  - Runtime code patching is allowed, via target-specific mechanisms.
  - However there is no guarantee that each block of assembly code in the source directly corresponds to a single instance of instructions in the object file; the compiler is free to duplicate or deduplicate the assembly code in asm! blocks.

r[asm.rules.stack-below-sp]

- Unless the nostack option is set, assembly code is allowed to use stack space below the stack pointer.
  - On entry to the assembly code the stack pointer is guaranteed to be suitably aligned (according to the target ABI) for a function call.
  - You are responsible for making sure you don't overflow the stack (e.g. use stack probing to ensure you hit a guard page).
  - You should adjust the stack pointer when allocating stack memory as required by the target ABI.
  - The stack pointer must be restored to its original value before leaving the assembly code.

r[asm.rules.noreturn]

• If the noreturn option is set then behavior is undefined if execution falls through the end of the assembly code.

r[asm.rules.pure]

• If the pure option is set then behavior is undefined if the asm! has sideeffects other than its direct outputs. Behavior is also undefined if two executions of the asm! code with the same inputs result in different outputs. • When used with the nomem option, "inputs" are just the direct inputs of the asm!.

• When used with the readonly option, "inputs" comprise the direct inputs of the assembly code and any memory that it is allowed to read. r[asm.rules.preserved-registers]

- These flags registers must be restored upon exiting the assembly code if the preserves\_flags option is set:
  - x86
    - Status flags in EFLAGS (CF, PF, AF, ZF, SF, OF).
    - Floating-point status word (all).
    - Floating-point exception flags in MXCSR (PE, UE, OE, ZE, DE, IE).
  - ARM
    - Condition flags in CPSR (N, Z, C, V)
    - Saturation flag in CPSR (Q)
    - Greater than or equal flags in CPSR (GE).
    - Condition flags in FPSCR (N, Z, C, V)
    - Saturation flag in FPSCR (QC)
    - Floating-point exception flags in FPSCR (IDC, IXC, UFC, OFC, DZC, IOC).
  - AArch64 and Arm64EC
    - Condition flags (NZCV register).
    - Floating-point status (FPSR register).
  - RISC-V
    - Floating-point exception flags in fcsr (fflags).
    - Vector extension state (vtype, vl, vcsr).
  - LoongArch
    - Floating-point condition flags in \$fcc[0-7].
  - s390x

• The condition code register cc.

r[asm.rules.x86-df]

- On x86, the direction flag (DF in EFLAGS) is clear on entry to the assembly code and must be clear on exit.
  - Behavior is undefined if the direction flag is set on exiting the assembly code.

r[asm.rules.x86-x87]

- On x86, the x87 floating-point register stack must remain unchanged unless all of the st([0-7]) registers have been marked as clobbered with out("st(0)") \_, out("st(1)") \_, .....
  - If all x87 registers are clobbered then the x87 register stack is guaranteed to be empty upon entering the assembly code. Assembly code must ensure that the x87 register stack is also empty when exiting the assembly code.

```
# #[cfg(target_arch = "x86_64")]
pub fn fadd(x: f64, y: f64) -> f64 {
  let mut out = 0f64;
  let mut top = 0u16;
  // we can do complex stuff with x87 if we clobber the entire x87
stack
  unsafe { core::arch::asm!(
    "fld qword ptr [{x}]",
    "fld qword ptr [{y}])",
    "faddp",
    "fstp qword ptr [{out}]",
    "xor eax, eax",
    "fstsw ax",
    "shl eax, 11",
    x = in(reg) \& x,
    y = in(reg) \& y,
    out = in(reg) &mut out,
    out("st(0)") _, out("st(1)") _, out("st(2)") _, out("st(3)") _,
    out("st(4)") _, out("st(5)") _, out("st(6)") _, out("st(7)") _,
    out("eax") top
```

```
);}
assert_eq!(top & 0x7, 0);
out
}
pub fn main() {
# #[cfg(target_arch = "x86_64")]{
    assert_eq!(fadd(1.0, 1.0), 2.0);
# }
}
```

r[asm.rules.arm64ec]

• On arm64ec, <u>call checkers with appropriate thunks</u> are mandatory when calling functions.

r[asm.rules.only-on-exit]

- The requirement of restoring the stack pointer and non-output registers to their original value only applies when exiting the assembly code.
  - This means that assembly code that does not fall through and does not jump to any label blocks, even if not marked noreturn, doesn't need to preserve these registers.
  - When returning to the assembly code of a different asm! block than you entered (e.g. for context switching), these registers must contain the value they had upon entering the asm! block that you are *exiting*.
    - You cannot exit the assembly code of an asm! block that has not been entered. Neither can you exit the assembly code of an asm! block whose assembly code has already been exited (without first entering it again).
    - You are responsible for switching any target-specific state (e.g. thread-local storage, stack bounds).
    - You cannot jump from an address in one asm! block to an address in another, even within the same function or block, without treating their contexts as potentially different and requiring context switching. You cannot assume that any particular value in those contexts (e.g. current stack pointer or temporary values below the

stack pointer) will remain unchanged between the two asm! blocks.

 The set of memory locations that you may access is the intersection of those allowed by the asm! blocks you entered and exited.

r[asm.rules.not-successive]

• You cannot assume that two asm! blocks adjacent in source code, even without any other code between them, will end up in successive addresses in the binary without any other instructions between them.

r[asm.rules.not-exactly-once]

• You cannot assume that an asm! block will appear exactly once in the output binary. The compiler is allowed to instantiate multiple copies of the asm! block, for example when the function containing it is inlined in multiple places.

r[asm.rules.x86-prefix-restriction]

On x86, inline assembly must not end with an instruction prefix (such as LOCK) that would apply to instructions generated by the compiler.

• The compiler is currently unable to detect this due to the way inline assembly is compiled, but may catch and reject this in the future. r[asm.rules.preserves\_flags]

[!NOTE] As a general rule, the flags covered by preserves\_flags are those which are *not* preserved when performing a function call.

r[asm.naked-rules]

### **Rules for naked inline assembly**

r[asm.naked-rules.intro] To avoid undefined behavior, these rules must be followed when using function-scope inline assembly in naked functions (naked\_asm!):

r[asm.naked-rules.reg-not-input]

- Any registers not used for function inputs according to the calling convention and function signature will contain an undefined value on entry to the naked\_asm! block.
  - An "undefined value" in the context of inline assembly means that the register can (non-deterministically) have any one of the possible values allowed by the architecture. Notably it is not the same as an LLVM undef which can have a different value every time you read it (since such a concept does not exist in assembly code).

r[asm.naked-rules.callee-saved-registers]

• All callee-saved registers must have the same value upon return as they had on entry.

r[asm.naked-rules.caller-saved-registers]

- Caller-saved registers may be used freely. r[asm.naked-rules.noreturn]
- Behavior is undefined if execution falls through past the end of the assembly code.
  - Every path through the assembly code is expected to terminate with a return instruction or to diverge.

r[asm.naked-rules.mem-same-as-ffi]

• The set of memory locations that assembly code is allowed to read and write are the same as those allowed for an FFI function.

r[asm.naked-rules.black-box]

• The compiler cannot assume that the instructions in the naked\_asm! block
 are the ones that will actually be executed.

- This effectively means that the compiler must treat the naked\_asm! as a
  black box and only take the interface specification into account, not the
  instructions themselves.
- Runtime code patching is allowed, via target-specific mechanisms.

r[asm.naked-rules.unwind]

• Unwinding out of a naked\_asm! block is allowed.

```
• For correct behavior, the appropriate assembler directives that emit unwinding metadata must be used.
```

```
# #[cfg(target_arch = "x86_64")] {
#[unsafe(naked)]
extern "sysv64-unwind" fn unwinding_naked() {
    core::arch::naked_asm!(
        // "CFI" here stands for "call frame information".
        ".cfi_startproc",
        // The CFA (canonical frame address) is the value of `rsp`
           // before the `call`, i.e. before the return address,
`rip`,
          // was pushed to `rsp`, so it's eight bytes higher in
memory
        // than `rsp` upon function entry (after `rip` has been
        // pushed).
        11
        // This is the default, so we don't have to write it.
        //".cfi_def_cfa rsp, 8",
        11
        // The traditional thing to do is to preserve the base
        // pointer, so we'll do that.
        "push rbp",
         // Since we've now extended the stack downward by 8 bytes
in
         // memory, we need to adjust the offset to the CFA from
`rsp`
        // by another 8 bytes.
        ".cfi_adjust_cfa_offset 8",
         // We also then annotate where we've stored the caller's
value
```

// of `rbp`, relative to the CFA, so that when unwinding into // the caller we can find it, in case we need it to calculate // the caller's CFA relative to it. 11 // Here, we've stored the caller's `rbp` starting 16 bytes // below the CFA. I.e., starting from the CFA, there's first // the `rip` (which starts 8 bytes below the CFA and continues // up to it), then there's the caller's `rbp` that we just // pushed. ".cfi\_offset rbp, -16", // As is traditional, we set the base pointer to the value of // the stack pointer. This way, the base pointer stays the // same throughout the function body. "mov rbp, rsp", // We can now track the offset to the CFA from the base // pointer. This means we don't need to make any further // adjustments until the end, as we don't change `rbp`. ".cfi\_def\_cfa\_register rbp", // We can now call a function that may panic. "call {f}", // Upon return, we restore `rbp` in preparation for returning // ourselves. "pop rbp", // Now that we've restored `rbp`, we must specify the offset // to the CFA again in terms of `rsp`. ".cfi\_def\_cfa rsp, 8", // Now we can return. "ret", ".cfi\_endproc", f = sym may\_panic, )

}

```
extern "sysv64-unwind" fn may_panic() {
    panic!("unwind");
}
# }
```

#### [!NOTE]

For more information on the cfi assembler directives above, see these resources:

- Using as CFI directives
- DWARF Debugging Information Format Version 5
- ImperialViolet CFI directives in assembly files

r[asm.validity]

## **Correctness and Validity**

r[asm.validity.necessary-but-not-sufficient] In addition to all of the previous rules, the string argument to asm! must ultimately become--- after all other arguments are evaluated, formatting is performed, and operands are translated---assembly that is both syntactically correct and semantically valid for the target architecture. The formatting rules allow the compiler to generate assembly with correct syntax. Rules concerning operands permit valid translation of Rust operands into and out of the assembly code. Adherence to these rules is necessary, but not sufficient, for the final expanded assembly to be both correct and valid. For instance:

- arguments may be placed in positions which are syntactically incorrect after formatting
- an instruction may be correctly written, but given architecturally invalid operands
- an architecturally unspecified instruction may be assembled into unspecified code
- a set of instructions, each correct and valid, may cause undefined behavior if placed in immediate succession

r[asm.validity.non-exhaustive] As a result, these rules are *non-exhaustive*. The compiler is not required to check the correctness and validity of the initial string nor the final assembly that is generated. The assembler may check for correctness

and validity but is not required to do so. When using asm!, a typographical error may be sufficient to make a program unsound, and the rules for assembly may include thousands of pages of architectural reference manuals. Programmers should exercise appropriate care, as invoking this unsafe capability comes with assuming the responsibility of not violating rules of both the compiler or the architecture.

r[asm.directives]

### **Directives Support**

r[asm.directives.subset-supported] Inline assembly supports a subset of the directives supported by both GNU AS and LLVM's internal assembler, given as follows. The result of using other directives is assembler-specific (and may cause an error, or may be accepted as-is).

r[asm.directives.stateful] If inline assembly includes any "stateful" directive that modifies how subsequent assembly is processed, the assembly code must undo the effects of any such directives before the inline assembly ends.

r[asm.directives.supported-directives] The following directives are guaranteed to be supported by the assembler:

- .2byte
- .4byte
- .8byte
- .align
- .alt\_entry
- .ascii
- .asciz
- .balign
- .balignl
- .balignw
- .bss
- .byte
- .comm
- .data
- .def
- .double
- .endef

- .equ •
- .equiv •
- .eqv •
- .fill
- .float
- .global
- .globl
- .inst
- .insn
- .lcomm
- .long
- .octa •
- .option
- .p2align
- .popsection
- .private\_extern
- .pushsection
- .quad
- .scl
- .section
- .set •
- .short
- .size
- .skip
- .sleb128
- .space
- .string
- .text
- .type
- .uleb128
- .word

```
# #[cfg(target_arch = "x86_64")] {
let bytes: *const u8;
let len: usize;
```

```
unsafe {
    core::arch::asm!(
        "jmp 3f", "2: .ascii \"Hello World!\"",
        "3: lea {bytes}, [2b+rip]",
        "mov {len}, 12",
        bytes = out(reg) bytes,
        len = out(reg) len
    );
}
let
                 s
                                 =
                                                unsafe
                                                                     {
core::str::from_utf8_unchecked(core::slice::from_raw_parts(bytes,
len)) };
assert_eq!(s, "Hello World!");
# }
```

r[asm.target-specific-directives]

### **Target Specific Directive Support**

r[asm.target-specific-directives.dwarf-unwinding]

### **Dwarf Unwinding**

The following directives are supported on ELF targets that support DWARF unwind info:

- .cfi\_adjust\_cfa\_offset
- .cfi\_def\_cfa
- .cfi\_def\_cfa\_offset
- .cfi\_def\_cfa\_register
- .cfi\_endproc
- .cfi\_escape
- .cfi\_lsda
- .cfi\_offset
- .cfi\_personality
- .cfi\_register
- .cfi\_rel\_offset
- .cfi\_remember\_state

- .cfi\_restore
- .cfi\_restore\_state
- .cfi\_return\_column
- .cfi\_same\_value
- .cfi\_sections
- .cfi\_signal\_frame
- .cfi\_startproc
- .cfi\_undefined
- .cfi\_window\_save

r[asm.target-specific-directives.structured-exception-handling]

#### **Structured Exception Handling**

On targets with structured exception Handling, the following additional directives are guaranteed to be supported:

- .seh\_endproc
- .seh\_endprologue
- .seh\_proc
- .seh\_pushreg
- .seh\_savereg
- .seh\_setframe
- .seh\_stackalloc

r[asm.target-specific-directives.x86]

#### x86 (32-bit and 64-bit)

On x86 targets, both 32-bit and 64-bit, the following additional directives are guaranteed to be supported:

- .nops
- .code16
- .code32
- .code64

Use of .code16, .code32, and .code64 directives are only supported if the state is reset to the default before exiting the assembly code. 32-bit x86 uses .code32 by default, and x86\_64 uses .code64 by default.

r[asm.target-specific-directives.arm-32-bit]

### ARM (32-bit)

On ARM, the following additional directives are guaranteed to be supported:

- .even
- .fnstart
- .fnend
- .save
- .movsp
- .code
- .thumb
- .thumb\_func

r[safety]

# Unsafety

r[safety.intro] Unsafe operations are those that can potentially violate the memory-safety guarantees of Rust's static semantics.

r[safety.unsafe-ops] The following language level features cannot be used in the safe subset of Rust:

r[safety.unsafe-deref]

- Dereferencing a <u>raw pointer</u>. r[safety.unsafe-static]
- Reading or writing a <u>mutable</u> or <u>external</u> static variable. r[safety.unsafe-union-access]
- Accessing a field of a <u>union</u>, other than to assign to it. r[safety.unsafe-call]
- Calling an unsafe function (including an intrinsic or foreign function). r[safety.unsafe-target-feature-call]
- Calling a safe function marked with a [target\_feature] [attributes.codegen.target\_feature] from a function that does not have a target\_feature attribute enabling the same features (see [attributes.codegen.target\_feature.safety-restrictions]).
   r[safety.unsafe-impl]
- Implementing an <u>unsafe trait</u>. r[safety.unsafe-extern]
- Declaring an <u>extern</u> block<sup>1</sup>.
   r[safety.unsafe-attribute]
- Applying an <u>unsafe attribute</u> to an item.

1

Prior to the 2024 edition, extern blocks were allowed to be declared without unsafe.

r[unsafe]

# The unsafe keyword

r[unsafe.intro] The unsafe keyword can occur in several different contexts: unsafe functions (unsafe fn), unsafe blocks (unsafe {}), unsafe traits (unsafe trait), unsafe trait implementations (unsafe impl), unsafe external blocks (unsafe extern), and unsafe attributes (# [unsafe(attr)]). It plays several different roles, depending on where it is used and whether the unsafe\_op\_in\_unsafe\_fn lint is enabled:

- it is used to mark code that *defines* extra safety conditions (unsafe
   fn, unsafe trait)
- it is used to mark code that needs to satisfy extra safety conditions (unsafe {}, unsafe impl, unsafe fn without unsafe op in unsafe fn, unsafe extern, #[unsafe(attr)])

The following discusses each of these cases. See the <u>keyword</u> <u>documentation</u> for some illustrative examples.

r[unsafe.fn]

### **Unsafe functions (unsafe fn)**

r[unsafe.fn.intro] Unsafe functions are functions that are not safe in all contexts and/or for all possible inputs. We say they have *extra safety conditions*, which are requirements that must be upheld by all callers and that the compiler does not check. For example, <u>get unchecked</u> has the extra safety condition that the index must be in-bounds. The unsafe function should come with documentation explaining what those extra safety conditions are.

r[unsafe.fn.safety] Such a function must be prefixed with the keyword unsafe and can only be called from inside an unsafe block, or inside unsafe fn without the <u>unsafe op in unsafe fn</u> lint.

r[unsafe.block]

# Unsafe blocks (unsafe {})

r[unsafe.block.intro] A block of code can be prefixed with the unsafe keyword to permit using the unsafe actions as defined in the <u>Unsafety</u> chapter, such as calling other unsafe functions or dereferencing raw pointers.

r[unsafe.block.fn-body] By default, the body of an unsafe function is also considered to be an unsafe block; this can be changed by enabling the <u>unsafe op in unsafe fn</u> lint.

By putting operations into an unsafe block, the programmer states that they have taken care of satisfying the extra safety conditions of all operations inside that block.

Unsafe blocks are the logical dual to unsafe functions: where unsafe functions define a proof obligation that callers must uphold, unsafe blocks state that all relevant proof obligations of functions or operations called inside the block have been discharged. There are many ways to discharge proof obligations; for example, there could be run-time checks or data structure invariants that guarantee that certain properties are definitely true, or the unsafe block could be inside an unsafe fn, in which case the block can use the proof obligations of that function to discharge the proof obligations arising inside the block.

Unsafe blocks are used to wrap foreign libraries, make direct use of hardware or implement features not directly present in the language. For example, Rust provides the language features necessary to implement memory-safe concurrency in the language but the implementation of threads and message passing in the standard library uses unsafe blocks.

Rust's type system is a conservative approximation of the dynamic safety requirements, so in some cases there is a performance cost to using safe code. For example, a doubly-linked list is not a tree structure and can only be represented with reference-counted pointers in safe code. By using unsafe blocks to represent the reverse links as raw pointers, it can be implemented without reference counting. (See <u>"Learn Rust With Entirely Too Many Linked Lists"</u> for a more in-depth exploration of this particular example.)

r[unsafe.trait]

## **Unsafe traits (unsafe trait)**

r[unsafe.trait.intro] An unsafe trait is a trait that comes with extra safety conditions that must be upheld by *implementations* of the trait. The unsafe trait should come with documentation explaining what those extra safety conditions are.

r[unsafe.trait.safety] Such a trait must be prefixed with the keyword unsafe and can only be implemented by unsafe impl blocks.

r[unsafe.impl]

# Unsafe trait implementations (unsafe impl)

When implementing an unsafe trait, the implementation needs to be prefixed with the unsafe keyword. By writing unsafe impl, the programmer states that they have taken care of satisfying the extra safety conditions required by the trait.

Unsafe trait implementations are the logical dual to unsafe traits: where unsafe traits define a proof obligation that implementations must uphold, unsafe implementations state that all relevant proof obligations have been discharged.

r[unsafe.extern]

# Unsafe external blocks (unsafe extern)

The programmer who declares an <u>external block</u> must assure that the signatures of the items contained within are correct. Failing to do so may lead to undefined behavior. That this obligation has been met is indicated by writing unsafe extern.

r[unsafe.extern.edition2024]

[!EDITION-2024] Prior to edition 2024, extern blocks were allowed without being qualified as unsafe.

r[unsafe.attribute]

## Unsafe attributes (#[unsafe(attr)])

An <u>unsafe attribute</u> is one that has extra safety conditions that must be upheld when using the attribute. The compiler cannot check whether these conditions have been upheld. To assert that they have been, these attributes must be wrapped in <code>unsafe(..)</code>, e.g. <code>#[unsafe(no\_mangle)]</code>.

r[undefined]

# **Behavior considered undefined**

r[undefined.general] Rust code is incorrect if it exhibits any of the behaviors in the following list. This includes code within unsafe blocks and unsafe functions. unsafe only means that avoiding undefined behavior is on the programmer; it does not change anything about the fact that Rust programs must never cause undefined behavior.

r[undefined.soundness] It is the programmer's responsibility when writing unsafe code to ensure that any safe code interacting with the unsafe code cannot trigger these behaviors. unsafe code that satisfies this property for any safe client is called *sound*; if unsafe code can be misused by safe code to exhibit undefined behavior, it is *unsound*.

[!WARNING] The following list is not exhaustive; it may grow or shrink. There is no formal model of Rust's semantics for what is and is not allowed in unsafe code, so there may be more behavior considered unsafe. We also reserve the right to make some of the behavior in that list defined in the future. In other words, this list does not say that anything will *definitely* always be undefined in all future Rust version (but we might make such commitments for some list items in the future).

Please read the <u>Rustonomicon</u> before writing unsafe code.

r[undefined.race]

• Data races.

r[undefined.pointer-access]

• Accessing (loading from or storing to) a place that is <u>dangling</u> or <u>based</u> <u>on a misaligned pointer</u>.

r[undefined.place-projection]

• Performing a place projection that violates the requirements of <u>in-bounds pointer arithmetic</u>. A place projection is a <u>field expression</u>, a <u>tuple index expression</u>, or an <u>array/slice index expression</u>.

r[undefined.alias]

- Breaking the pointer aliasing rules. The exact aliasing rules are not determined yet, but here is an outline of the general principles: &T must point to memory that is not mutated while they are live (except for data inside an UnsafeCell<U>), and &mut T must point to memory that is not read or written by any pointer not derived from the reference and that no other reference points to while they are live.
   Box<T> is treated similar to &'static mut T for the purpose of these rules. The exact liveness duration is not specified, but some bounds exist:
  - For references, the liveness duration is upper-bounded by the syntactic lifetime assigned by the borrow checker; it cannot be live any *longer* than that lifetime.
  - Each time a reference or box is dereferenced or reborrowed, it is considered live.
  - Each time a reference or box is passed to or returned from a function, it is considered live.
  - When a reference (but not a Box !) is passed to a function, it is live at least as long as that function call, again except if the &T contains an <u>UnsafeCell<U></u>.

All this also applies when values of these types are passed in a (nested) field of a compound type, but not behind pointer indirections. r[undefined.immutable]

Mutating immutable bytes. All bytes reachable through a <u>const-promoted</u> expression are immutable, as well as bytes reachable through borrows in <u>static</u> and <u>const</u> initializers that have been <u>lifetime-extended</u> to <u>'static</u>. The bytes owned by an immutable binding or immutable <u>static</u> are immutable, unless those bytes are part of an <u>UnsafeCell<U></u>.

Moreover, the bytes <u>pointed to</u> by a shared reference, including transitively through other references (both shared and mutable) and

Box es, are immutable; transitivity includes those references stored in fields of compound types.

A mutation is any write of more than 0 bytes which overlaps with any of the relevant bytes (even if that write does not change the memory contents).

r[undefined.intrinsic]

• Invoking undefined behavior via compiler intrinsics. r[undefined.target-feature]

• Executing code compiled with platform features that the current platform does not support (see <u>target feature</u>), *except* if the platform explicitly documents this to be safe.

r[undefined.call]

 Calling a function with the wrong <u>call ABI</u>, or unwinding past a stack frame that does not allow unwinding (e.g. by calling a "C-unwind" function imported or transmuted as a "C" function or function pointer).

r[undefined.invalid]

• Producing an <u>invalid value</u>. "Producing" a value happens any time a value is assigned to or read from a place, passed to a function/primitive operation or returned from a function/primitive operation.

r[undefined.asm]

• Incorrect use of inline assembly. For more details, refer to the <u>rules</u> to follow when writing code that uses inline assembly.

r[undefined.const-transmute-ptr2int]

• In <u>const context</u>: transmuting or otherwise reinterpreting a pointer (reference, raw pointer, or function pointer) into some allocated object as a non-pointer type (such as integers). 'Reinterpreting' refers to loading the pointer value at integer type without a cast, e.g. by doing raw pointer casts or using a union.
r[undefined.runtime]

- Violating assumptions of the Rust runtime. Most assumptions of the Rust runtime are currently not explicitly documented.
  - For assumptions specifically related to unwinding, see the <u>panic</u> <u>documentation</u>.
  - The runtime assumes that a Rust stack frame is not deallocated without executing destructors for local variables owned by the stack frame. This assumption can be violated by C functions like longjmp.

[!NOTE] Undefined behavior affects the entire program. For example, calling a function in C that exhibits undefined behavior of C means your entire program contains undefined behaviour that can also affect the Rust code. And vice versa, undefined behavior in Rust can cause adverse affects on code executed by any FFI calls to other languages.

r[undefined.pointed-to]

#### **Pointed-to bytes**

The span of bytes a pointer or reference "points to" is determined by the pointer value and the size of the pointee type (using size\_of\_val).

r[undefined.misaligned]

#### **Places based on misaligned pointers**

r[undefined.misaligned.general] A place is said to be "based on a misaligned pointer" if the last \* projection during place computation was performed on a pointer that was not aligned for its type. (If there is no \* projection in the place expression, then this is accessing the field of a local or static and rustc will guarantee proper alignment. If there are multiple \* projection, then each of them incurs a load of the pointer-to-be-dereferenced itself from memory, and each of these loads is subject to the alignment constraint. Note that some \* projections can be omitted in surface Rust syntax due to automatic dereferencing; we are considering the fully expanded place expression here.)

For instance, if ptr has type \*const S where S has an alignment of 8, then ptr must be 8-aligned or else (\*ptr).f is "based on an misaligned pointer". This is true even if the type of the field f is u8 (i.e., a type with alignment 1). In other words, the alignment requirement derives from the type of the pointer that was dereferenced, *not* the type of the field that is being accessed.

r[undefined.misaligned.load-store] Note that a place based on a misaligned pointer only leads to Undefined Behavior when it is loaded from or stored to.

r[undefined.misaligned.raw] &raw const/&raw mut on such a place is allowed.

r[undefined.misaligned.reference] &/&mut on a place requires the alignment of the field type (or else the program would be "producing an invalid value"), which generally is a less restrictive requirement than being based on an aligned pointer.

r[undefined.misaligned.packed] Taking a reference will lead to a compiler error in cases where the field type might be more aligned than the type that contains it, i.e., repr(packed). This means that being based on an aligned pointer is always sufficient to ensure that the new reference is aligned, but it is not always necessary.

r[undefined.dangling]

#### **Dangling pointers**

r[undefined.dangling.general] A reference/pointer is "dangling" if not all of the bytes it <u>points to</u> are part of the same live allocation (so in particular they all have to be part of *some* allocation).

r[undefined.dangling.zero-size] If the size is 0, then the pointer is trivially never "dangling" (even if it is a null pointer).

r[undefined.dangling.dynamic-size] Note that dynamically sized types (such as slices and strings) point to their entire range, so it is important that the length metadata is never too large.

r[undefined.dangling.alloc-limit] In particular, the dynamic size of a Rust value (as determined by size\_of\_val) must never exceed isize::MAX, since it is impossible for a single allocation to be larger than isize::MAX.

r[undefined.validity]

#### **Invalid values**

r[undefined.validity.general] The Rust compiler assumes that all values produced during program execution are "valid", and producing an invalid value is hence immediate UB.

Whether a value is valid depends on the type: r[undefined.validity.bool]

- A <u>bool</u> value must be false (0) or true (1). r[undefined.validity.fn-pointer]
- A fn pointer value must be non-null. r[undefined.validity.char]
- A char value must not be a surrogate (i.e., must not be in the range 0xD800..=0xDFFF) and must be equal to or less than char::MAX.
   r[undefined.validity.never]

• A ! value must never exist.

r[undefined.validity.scalar]

- An integer (i\*/u\*), floating point value (f\*), or raw pointer must be initialized, i.e., must not be obtained from uninitialized memory.
   r[undefined.validity.str]
- A str value is treated like [u8], i.e. it must be initialized. r[undefined.validity.enum]
- An enum must have a valid discriminant, and all fields of the variant indicated by that discriminant must be valid at their respective type. r[undefined.validity.struct]
- A struct, tuple, and array requires all fields/elements to be valid at their respective type.

r[undefined.validity.union]

For a union, the exact validity requirements are not decided yet.
 Obviously, all values that can be created entirely in safe code are valid.
 If the union has a zero-sized field, then every possible value is valid.
 Further details are <u>still being debated</u>.

r[undefined.validity.reference-box]

A reference or [Box<T>] must be aligned and non-null, it cannot be dangling, and it must point to a valid value (in case of dynamically sized types, using the actual dynamic type of the pointee as determined by the metadata). Note that the last point (about pointing to a valid value) remains a subject of some debate.

r[undefined.validity.wide]

- The metadata of a wide reference, [Box<T>], or raw pointer must match the type of the unsized tail:
  - dyn Trait metadata must be a pointer to a compiler-generated vtable for Trait. (For raw pointers, this requirement remains a subject of some debate.)
  - Slice ([T]) metadata must be a valid usize. Furthermore, for wide references and [Box<T>], slice metadata is invalid if it makes the total size of the pointed-to value bigger than isize::MAX.

r[undefined.validity.valid-range]

 If a type has a custom range of a valid values, then a valid value must be in that range. In the standard library, this affects <u>NonNull<T></u> and <u>NonZero<T></u>.

```
[!NOTE] rustc achieves this with the unstable
rustc_layout_scalar_valid_range_* attributes.
```

r[undefined.validity.undef] **Note:** Uninitialized memory is also implicitly invalid for any type that has a restricted set of valid values. In

other words, the only cases in which reading uninitialized memory is permitted are inside unions and in "padding" (the gaps between the fields of a type).

## **Behavior not considered unsafe**

The Rust compiler does not consider the following behaviors *unsafe*, though a programmer may (should) find them undesirable, unexpected, or erroneous.

- Deadlocks
- Leaks of memory and other resources
- Exiting without calling destructors
- Exposing randomized base addresses through pointer leaks

#### **Integer overflow**

If a program contains arithmetic overflow, the programmer has made an error. In the following discussion, we maintain a distinction between arithmetic overflow and wrapping arithmetic. The first is erroneous, while the second is intentional.

When the programmer has enabled debug\_assert! assertions (for example, by enabling a non-optimized build), implementations must insert dynamic checks that panic on overflow. Other kinds of builds may result in panics or silently wrapped values on overflow, at the implementation's discretion.

In the case of implicitly-wrapped overflow, implementations must provide well-defined (even if still considered erroneous) results by using two's complement overflow conventions.

The integral types provide inherent methods to allow programmers explicitly to perform wrapping arithmetic. For example, i32::wrapping\_add provides two's complement, wrapping addition.

The standard library also provides a Wrapping<T> newtype which ensures all standard arithmetic operations for T have wrapping semantics.

See <u>RFC 560</u> for error conditions, rationale, and more details about integer overflow.

#### **Logic errors**

Safe code may impose extra logical constraints that can be checked at neither compile-time nor runtime. If a program breaks such a constraint, the behavior may be unspecified but will not result in undefined behavior. This could include panics, incorrect results, aborts, and non-termination. The behavior may also differ between runs, builds, or kinds of build.

For example, implementing both Hash and Eq requires that values considered equal have equal hashes. Another example are data structures like BinaryHeap, BTreeMap, BTreeSet, HashMap and HashSet which describe constraints on the modification of their keys while they are in the data structure. Violating such constraints is not considered unsafe, yet the program is considered erroneous and its behavior unpredictable.

r[const-eval]

### **Constant evaluation**

r[const-eval.general] Constant evaluation is the process of computing the result of <u>expressions</u> during compilation. Only a subset of all expressions can be evaluated at compile-time.

```
r[const-eval.const-expr]
```

#### **Constant expressions**

r[const-eval.const-expr.general] Certain forms of expressions, called constant expressions, can be evaluated at compile time.

r[const-eval.const-expr.const-context] In <u>const contexts</u>, these are the only allowed expressions, and are always evaluated at compile time.

r[const-eval.const-expr.runtime-context] In other places, such as <u>let</u> <u>statements</u>, constant expressions *may* be, but are not guaranteed to be, evaluated at compile time.

r[const-eval.const-expr.error] Behaviors such as out of bounds <u>array</u> <u>indexing</u> or <u>overflow</u> are compiler errors if the value must be evaluated at compile time (i.e. in const contexts). Otherwise, these behaviors are warnings, but will likely panic at run-time.

r[const-eval.const-expr.list] The following expressions are constant expressions, so long as any operands are also constant expressions and do not cause any <u>Drop::drop</u> calls to be run.

r[const-eval.const-expr.literal]

```
• <u>Literals</u>.
```

```
r[const-eval.const-expr.parameter]
```

```
• <u>Const parameters</u>.
```

```
r[const-eval.const-expr.path-item]
```

• <u>Paths</u> to <u>functions</u> and <u>constants</u>. Recursively defining constants is not allowed.

r[const-eval.const-expr.path-static]

- Paths to <u>statics</u> with these restrictions:
  - Writes to static items are not allowed in any constant evaluation context.
  - Reads from extern statics are not allowed in any constant evaluation context.

 If the evaluation is *not* carried out in an initializer of a static item, then reads from any mutable static are not allowed. A mutable static is a static mut item, or a static item with an interior-mutable type.

These requirements are checked only when the constant is evaluated. In other words, having such accesses syntactically occur in const contexts is allowed as long as they never get executed.

r[const-eval.const-expr.tuple]

• <u>Tuple expressions</u>. r[const-eval.const-expr.array]

• <u>Array expressions</u>.

r[const-eval.const-expr.constructor]

• <u>Struct</u> expressions. r[const-eval.const-expr.block]

- <u>Block expressions</u>, including unsafe and const blocks.
  - <u>let statements</u> and thus irrefutable <u>patterns</u>, including mutable bindings
  - <u>assignment expressions</u>
  - <u>compound assignment expressions</u>
  - <u>expression statements</u>

r[const-eval.const-expr.field]

• <u>Field</u> expressions. r[const-eval.const-expr.index]

• Index expressions, <u>array indexing</u> or <u>slice</u> with a usize. r[const-eval.const-expr.range]

• <u>Range expressions</u>. r[const-eval.const-expr.closure] • <u>Closure expressions</u> which don't capture variables from the environment.

r[const-eval.const-expr.builtin-arith-logic]

• Built-in <u>negation</u>, <u>arithmetic</u>, <u>logical</u>, <u>comparison</u> or <u>lazy boolean</u> operators used on integer and floating point types, bool, and char.

r[const-eval.const-expr.borrows]

- All forms of <u>borrow</u>s, including raw borrows, except borrows of expressions whose temporary scopes would be extended (see <u>temporary lifetime extension</u>) to the end of the program and which are either:
  - Mutable borrows.
  - Shared borrows of expressions that result in values with <u>interior</u> <u>mutability</u>.

```
// Due to being in tail position, this borrow extends the
scope of the
// temporary to the end of the program. Since the borrow is
mutable,
// this is not allowed in a const expression.
const C: &u8 = &mut 0; // ERROR not allowed
// Const blocks are similar to initializers of `const`
items.
let _: &u8 = const { &mut 0 }; // ERROR not allowed
# use core::sync::atomic::AtomicU8;
// This is not allowed as 1) the temporary scope is
extended to the
// end of the program and 2) the temporary has interior
mutability.
const C: &AtomicU8 = &AtomicU8::new(0); // ERROR not
allowed
# use core::sync::atomic::AtomicU8;
// As above.
let _: &_ = const { &AtomicU8::new(0) }; // ERROR not
allowed
```

```
# #![allow(static_mut_refs)]
// Even though this borrow is mutable, it's not of a
temporary, so
// this is allowed.
const C: &u8 = unsafe { static mut S: u8 = 0; &mut S };
// OK
# use core::sync::atomic::AtomicU8;
// Even though this borrow is of a value with interior
mutability,
// it's not of a temporary, so this is allowed.
const C: &AtomicU8 = {
    static S: AtomicU8 = AtomicU8::new(0); &S // OK
};
# use core::sync::atomic::AtomicU8;
// This shared borrow of an interior mutable temporary is
allowed
// because its scope is not extended.
const C: () = { _ = &AtomicU8::new(0); }; // OK
// Even though the borrow is mutable and the temporary
lives to the
// end of the program due to promotion, this is allowed
because the
// borrow is not in tail position and so the scope of the
temporary
// is not extended via temporary lifetime extension.
const C: () = { let _: &'static mut [u8] = &mut []; }; //
0K
11
11
                                       Promoted
temporary.
```

[!NOTE] In other words --- to focus on what's allowed rather than what's not allowed --- shared borrows of interior mutable data and mutable borrows are only allowed in a <u>const context</u> when the borrowed <u>place expression</u> is *transient*, *indirect*, or *static*.

A place expression is *transient* if it is a variable local to the current const context or an expression whose temporary scope is contained inside the current const context.

```
// The borrow is of a variable local to the
initializer, therefore
// this place expression is transient.
const C: () = { let mut x = 0; _ = &mut x; };
// The borrow is of a temporary whose scope has not
been extended,
// therefore this place expression is transient.
const C: () = { _ = &mut Ou8; };
// When a temporary is promoted but not lifetime
extended, its
// place expression is still treated as transient.
const C: () = { let _: &'static mut [u8] = &mut [];
};
```

```
A place expression is indirect if it is a <u>dereference expression</u>.
const C: () = { _ = &mut *(&mut 0); };
```

A place expression is *static* if it is a **static** item.

```
# #![allow(static_mut_refs)]
const C: &u8 = unsafe { static mut S: u8 = 0; &mut S
};
```

[!NOTE] One surprising consequence of these rules is that we allow this,

```
const C: &[u8] = { let x: &mut [u8] = &mut []; x };
// OK
// ~~~~~~
// Empty arrays are promoted even behind mutable
borrows.
```

but we disallow this similar code:

The difference between these is that, in the first, the empty array is <u>promoted</u> but its scope does not undergo <u>temporary</u> <u>lifetime extension</u>, so we consider the <u>place expression</u> to be transient (even though after promotion the place indeed lives to the end of the program). In the second, the scope of the empty array temporary does undergo lifetime extension, and so it is rejected due to being a mutable borrow of a lifetime-extended temporary (and therefore borrowing a non-transient place expression).

The effect is surprising because temporary lifetime extension, in this case, causes less code to compile than would without it.

See <u>issue #143129</u> for more details.

r[const-eval.const-expr.deref]

• The <u>dereference operator</u> except for raw pointers. r[const-eval.const-expr.group]

- <u>Grouped</u> expressions. r[const-eval.const-expr.cast]
- <u>Cast</u> expressions, except
  - pointer to address casts and
  - function pointer to address casts.

r[const-eval.const-expr.const-fn]

• Calls of <u>const functions</u> and const methods. r[const-eval.const-expr.loop]

• <u>loop</u> and <u>while</u> expressions. r[const-eval.const-expr.if-match] • <u>if</u> and <u>match</u> expressions. r[const-eval.const-context]

#### **Const context**

r[const-eval.const-context.general] A *const context* is one of the following:

r[const-eval.const-context.array-length]

- <u>Array type length expressions</u> r[const-eval.const-context.repeat-length]
- <u>Array repeat length expressions</u> r[const-eval.const-context.init]
- The initializer of
  - <u>constants</u>
  - <u>statics</u>
  - enum discriminants

r[const-eval.const-context.generic]

• A <u>const generic argument</u> r[const-eval.const-context.block]

• A <u>const block</u>

Const contexts that are used as parts of types (array type and repeat length expressions as well as const generic arguments) can only make restricted use of surrounding generic parameters: such an expression must either be a single bare const generic parameter, or an arbitrary expression not making use of any generics.

r[const-eval.const-fn]

#### **Const Functions**

r[const-eval.const-fn.general] A *const fn* is a function that one is permitted to call from a const context.

r[const-eval.const-fn.usage] Declaring a function **const** has no effect on any existing uses, it only restricts the types that arguments and the return type may use, and restricts the function body to constant expressions.

r[const-eval.const-fn.const-context] When called from a const context, the function is interpreted by the compiler at compile time. The interpretation happens in the environment of the compilation target and not the host. So usize is 32 bits if you are compiling against a 32 bit system, irrelevant of whether you are building on a 64 bit or a 32 bit system.

r[abi]

## **Application Binary Interface (ABI)**

r[abi.intro] This section documents features that affect the ABI of the compiled output of a crate.

See <u>extern functions</u> for information on specifying the ABI for exporting functions. See <u>external blocks</u> for information on specifying the ABI for linking external libraries.

r[abi.used]

#### The used attribute

r[abi.used.intro] The *used attribute* can only be applied to <u>static</u> <u>items</u>. This <u>attribute</u> forces the compiler to keep the variable in the output object file (.o, .rlib, etc. excluding final binaries) even if the variable is not used, or referenced, by any other item in the crate. However, the linker is still free to remove such an item.

Below is an example that shows under what conditions the compiler keeps a static item in the output object file.

```
// foo.rs
// This is kept because of `#[used]`:
#[used]
static F00: u32 = 0;
// This is removable because it is unused:
#[allow(dead_code)]
static BAR: u32 = 0;
// This is kept because it is publicly reachable:
pub static BAZ: u32 = 0;
11
   This is kept because it is referenced by a public,
reachable function:
static QUUX: u32 = 0;
pub fn quux() -> &'static u32 {
    &QUUX
}
// This is removable because it is referenced by a private,
unused (dead) function:
static CORGE: u32 = 0;
```

```
#[allow(dead_code)]
fn corge() -> &'static u32 {
    &CORGE
}
$ rustc -0 --emit=obj --crate-type=rlib foo.rs
$ nm -C foo.0
000000000000000 R foo::BAZ
00000000000000 r foo::FO0
0000000000000 R foo::QUUX
000000000000 T foo::quux
r[abi.no_mangle]
```

#### The no\_mangle attribute

r[abi.no\_mangle.intro] The *no\_mangle attribute* may be used on any <u>item</u> to disable standard symbol name mangling. The symbol for the item will be the identifier of the item's name.

r[abi.no\_mangle.publicly-exported] Additionally, the item will be publicly exported from the produced library or object file, similar to the <u>used attribute</u>.

r[abi.no\_mangle.unsafe] This attribute is unsafe as an unmangled symbol may collide with another symbol with the same name (or with a well-known symbol), leading to undefined behavior.

```
#[unsafe(no_mangle)]
extern "C" fn foo() {}
```

r[abi.no\_mangle.edition2024]

[!EDITION-2024] Before the 2024 edition it is allowed to use the no\_mangle attribute without the unsafe qualification.

r[abi.link\_section]

#### The link\_section attribute

r[abi.link\_section.intro] The *link\_section attribute* specifies the section of the object file that a <u>function</u> or <u>static</u>'s content will be placed into.

r[abi.link\_section.syntax] The link\_section attribute uses the
[MetaNameValueStr] syntax to specify the section name.

#[unsafe(no\_mangle)]

```
#[unsafe(link_section = ".example_section")]
```

```
pub static VAR1: u32 = 1;
```

r[abi.link\_section.unsafe] This attribute is unsafe as it allows users to place data and code into sections of memory not expecting them, such as mutable data into read-only areas.

r[abi.link\_section.edition2024]

[!EDITION-2024] Before the 2024 edition it is allowed to use the link\_section attribute without the unsafe qualification.

```
r[abi.export_name]
```

#### The export\_name attribute

r[abi.export\_name.intro] The *export\_name attribute* specifies the name of the symbol that will be exported on a <u>function</u> or <u>static</u>.

r[abi.export\_name.syntax] The export\_name attribute uses the [MetaNameValueStr] syntax to specify the symbol name.

```
#[unsafe(export_name = "exported_symbol_name")]
pub fn name_in_rust() { }
```

r[abi.export\_name.unsafe] This attribute is unsafe as a symbol with a custom name may collide with another symbol with the same name (or with a well-known symbol), leading to undefined behavior.

r[abi.export\_name.edition2024]

[!EDITION-2024] Before the 2024 edition it is allowed to use the export\_name attribute without the unsafe qualification.

r[runtime]

### The Rust runtime

This section documents features that define some aspects of the Rust runtime.

r[runtime.global\_allocator]

#### The global\_allocator attribute

The *global\_allocator attribute* is used on a <u>static item</u> implementing the <u>GlobalAlloc</u> trait to set the global allocator.

r[runtime.windows\_subsystem]

#### The windows\_subsystem attribute

r[runtime.windows\_subsystem.intro] The *windows\_subsystem* attribute may be applied at the crate level to set the <u>subsystem</u> when linking on a Windows target.

r[runtime.windows\_subsystem.syntax] It uses the [MetaNameValueStr] syntax to specify the subsystem with a value of either console or windows.

r[runtime.windows\_subsystem.ignored] This attribute is ignored on non-Windows targets, and for non-bin <u>crate types</u>.

r[runtime.windows\_subsystem.console] The "console" subsystem is the default. If a console process is run from an existing console then it will be attached to that console, otherwise a new console window will be created.

r[runtime.windows\_subsystem.windows] The "windows" subsystem is commonly used by GUI applications that do not want to display a console window on startup. It will run detached from any existing console.

```
#![windows_subsystem = "windows"]
```

## Appendices

### **Grammar summary**

The following is a summary of the grammar production rules. {{ grammar-summary }}

r[macro.ambiguity]

# Appendix: Macro Follow-Set Ambiguity Formal Specification

This page documents the formal specification of the follow rules for <u>Macros By Example</u>. They were originally specified in <u>RFC 550</u>, from which the bulk of this text is copied, and expanded upon in subsequent RFCs.

r[macro.ambiguity.convention]
# **Definitions & Conventions**

r[macro.ambiguity.convention.defs]

- macro: anything invocable as foo!(...) in source code.
- MBE: macro-by-example, a macro defined by macro\_rules.
- matcher: the left-hand-side of a rule in a macro\_rules invocation, or a subportion thereof.
- macro parser: the bit of code in the Rust parser that will parse the input using a grammar derived from all of the matchers.
- fragment: The class of Rust syntax that a given matcher will accept (or "match").
- repetition : a fragment that follows a regular repeating pattern
- NT: non-terminal, the various "meta-variables" or repetition matchers that can appear in a matcher, specified in MBE syntax with a leading \$ character.
- **simple NT**: a "meta-variable" non-terminal (further discussion below).
- complex NT: a repetition matching non-terminal, specified via repetition operators (\*, +, ?).
- token: an atomic element of a matcher; i.e. identifiers, operators, open/close delimiters, *and* simple NT's.
- token tree: a tree structure formed from tokens (the leaves), complex NT's, and finite sequences of token trees.
- delimiter token: a token that is meant to divide the end of one fragment and the start of the next fragment.
- separator token: an optional delimiter token in an complex NT that separates each pair of elements in the matched repetition.
- separated complex NT: a complex NT that has its own separator token.
- delimited sequence: a sequence of token trees with appropriate open- and close-delimiters at the start and end of the sequence.

- empty fragment: The class of invisible Rust syntax that separates tokens, i.e. whitespace, or (in some lexical contexts), the empty token sequence.
- fragment specifier: The identifier in a simple NT that specifies which fragment the NT accepts.
- language: a context-free language.

Example:

```
macro_rules! i_am_an_mbe {
    (start $foo:expr $($i:ident),* end) => ($foo)
}
```

r[macro.ambiguity.convention.matcher] (start \$foo:expr
\$(\$i:ident),\* end) is a matcher. The whole matcher is a delimited
sequence (with open- and close-delimiters ( and )), and \$foo and \$i are
simple NT's with expr and ident as their respective fragment specifiers.

r[macro.ambiguity.convention.complex-nt] \$(i:ident),\* is also an
NT; it is a complex NT that matches a comma-separated repetition of
identifiers. The , is the separator token for the complex NT; it occurs in
between each pair of elements (if any) of the matched fragment.

Another example of a complex NT is \$(hi \$e:expr ;)+, which matches any fragment of the form hi <expr>; hi <expr>; hi <expr>; ... where hi <expr>; occurs at least once. Note that this complex NT does not have a dedicated separator token.

(Note that Rust's parser ensures that delimited sequences always occur with proper nesting of token tree structure and correct matching of openand close-delimiters.)

r[macro.ambiguity.convention.vars] We will tend to use the variable "M" to stand for a matcher, variables "t" and "u" for arbitrary individual tokens, and the variables "tt" and "uu" for arbitrary token trees. (The use of "tt" does present potential ambiguity with its additional role as a fragment specifier; but it will be clear from context which interpretation is meant.)

r[macro.ambiguity.convention.set] "SEP" will range over separator tokens, "OP" over the repetition operators **\***, +, and ?, "OPEN"/"CLOSE"

over matching token pairs surrounding a delimited sequence (e.g. [ and ]).

r[macro.ambiguity.convention.sequence-vars] Greek letters " $\alpha$ " " $\beta$ " " $\gamma$ " " $\delta$ " stand for potentially empty token-tree sequences. (However, the Greek letter " $\epsilon$ " (epsilon) has a special role in the presentation and does not stand for a token-tree sequence.)

• This Greek letter convention is usually just employed when the presence of a sequence is a technical detail; in particular, when we wish to *emphasize* that we are operating on a sequence of token-trees, we will use the notation "tt ..." for the sequence, not a Greek letter.

Note that a matcher is merely a token tree. A "simple NT", as mentioned above, is an meta-variable NT; thus it is a non-repetition. For example, \$foo:ty is a simple NT but \$(\$foo:ty)+ is a complex NT.

Note also that in the context of this formalism, the term "token" generally *includes* simple NTs.

Finally, it is useful for the reader to keep in mind that according to the definitions of this formalism, no simple NT matches the empty fragment, and likewise no token matches the empty fragment of Rust syntax. (Thus, the *only* NT that can match the empty fragment is a complex NT.) This is not actually true, because the vis matcher can match an empty fragment. Thus, for the purposes of the formalism, we will treat \$v:vis as actually being \$(\$v:vis)?, with a requirement that the matcher match an empty fragment.

r[macro.ambiguity.invariant]

#### The Matcher Invariants

r[macro.ambiguity.invariant.list] To be valid, a matcher must meet the following three invariants. The definitions of FIRST and FOLLOW are described later.

```
1. For any two successive token tree sequences in a matcher M (i.e. M =
```

```
... tt uu ...) with uu ... nonempty, we must have FOLLOW(... tt) U {\epsilon} \supseteq FIRST(uu ...).
```

- 2. For any separated complex NT in a matcher, M = ... \$(tt ...) SEP
  OP ..., we must have SEP ∈ FOLLOW(tt ...).
- 3. For an unseparated complex NT in a matcher, M = ... \$(tt ...) OP
   ..., if OP = \* or +, we must have FOLLOW(tt ...) ⊇ FIRST(tt
   ...).

r[macro.ambiguity.invariant.follow-matcher] The first invariant says that whatever actual token that comes after a matcher, if any, must be somewhere in the predetermined follow set. This ensures that a legal macro definition will continue to assign the same determination as to where ... tt ends and uu ... begins, even as new syntactic forms are added to the language.

r[macro.ambiguity.invariant.separated-complex-nt] The second invariant says that a separated complex NT must use a separator token that is part of the predetermined follow set for the internal contents of the NT. This ensures that a legal macro definition will continue to parse an input fragment into the same delimited sequence of tt ....'s, even as new syntactic forms are added to the language.

r[macro.ambiguity.invariant.unseparated-complex-nt] The third invariant says that when we have a complex NT that can match two or more copies of the same thing with no separation in between, it must be permissible for them to be placed next to each other as per the first invariant. This invariant also requires they be nonempty, which eliminates a possible ambiguity.

NOTE: The third invariant is currently unenforced due to historical oversight and significant reliance on the behaviour. It is currently undecided what to do about this going forward. Macros that do not respect the behaviour may become invalid in a future edition of Rust. See the <u>tracking issue</u>.

r[macro.ambiguity.sets]

#### FIRST and FOLLOW, informally

r[macro.ambiguity.sets.intro] A given matcher M maps to three sets: FIRST(M), LAST(M) and FOLLOW(M).

Each of the three sets is made up of tokens. FIRST(M) and LAST(M) may also contain a distinguished non-token element  $\epsilon$  ("epsilon"), which

indicates that M can match the empty fragment. (But FOLLOW(M) is always just a set of tokens.)

Informally:

r[macro.ambiguity.sets.first]

• FIRST(M): collects the tokens potentially used first when matching a fragment to M.

r[macro.ambiguity.sets.last]

• LAST(M): collects the tokens potentially used last when matching a fragment to M.

r[macro.ambiguity.sets.follow]

• FOLLOW(M): the set of tokens allowed to follow immediately after some fragment matched by M.

In other words:  $t \in FOLLOW(M)$  if and only if there exists (potentially empty) token sequences  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\delta$  where:

- M matches  $\beta$ ,
- t matches γ, and
- The concatenation  $\alpha \beta \gamma \delta$  is a parseable Rust program.

r[macro.ambiguity.sets.universe] We use the shorthand ANYTOKEN to denote the set of all tokens (including simple NTs). For example, if any token is legal after a matcher M, then FOLLOW(M) = ANYTOKEN.

(To review one's understanding of the above informal descriptions, the reader at this point may want to jump ahead to the <u>examples of FIRST/LAST</u> before reading their formal definitions.)

r[macro.ambiguity.sets.def]

#### FIRST, LAST

r[macro.ambiguity.sets.def.intro] Below are formal inductive definitions for FIRST and LAST.

r[macro.ambiguity.sets.def.notation] "A  $\cup$  B" denotes set union, "A  $\cap$  B" denotes set intersection, and "A  $\setminus$  B" denotes set difference (i.e. all

elements of A that are not present in B).

r[macro.ambiguity.sets.def.first]

#### FIRST

r[macro.ambiguity.sets.def.first.intro] FIRST(M) is defined by case analysis on the sequence M and the structure of its first token-tree (if any):

r[macro.ambiguity.sets.def.first.epsilon]

 if M is the empty sequence, then FIRST(M) = { ε }, r[macro.ambiguity.sets.def.first.token]

 if M starts with a token t, then FIRST(M) = { t }, (Note: this covers the case where M starts with a delimited tokentree sequence, M = OPEN tt ... CLOSE ..., in which case t = OPEN and thus FIRST(M) = { OPEN }.)

(Note: this critically relies on the property that no simple NT matches the empty fragment.)

r[macro.ambiguity.sets.def.first.complex]

- Otherwise, M is a token-tree sequence starting with a complex NT:
   M = \$( tt ... ) OP α, or M = \$( tt ... ) SEP OP α, (where α is the (potentially empty) sequence of token trees for the rest of the matcher).
  - Let SEP\_SET(M) = { SEP } if SEP is present and ε ∈ FIRST(tt
     ...); otherwise SEP\_SET(M) = {}.
- Let ALPHA\_SET(M) = FIRST(α) if OP = \* or ? and ALPHA\_SET(M) = {} if OP = +.
- $FIRST(M) = (FIRST(tt ...) \setminus \{\epsilon\}) \cup SEP_SET(M) \cup ALPHA_SET(M).$

The definition for complex NTs deserves some justification. SEP\_SET(M) defines the possibility that the separator could be a valid first token for M, which happens when there is a separator defined and the repeated fragment could be empty. ALPHA\_SET(M) defines the possibility

that the complex NT could be empty, meaning that M's valid first tokens are those of the following token-tree sequences  $\alpha$ . This occurs when either \* or ? is used, in which case there could be zero repetitions. In theory, this could also occur if + was used with a potentially-empty repeating fragment, but this is forbidden by the third invariant.

From there, clearly FIRST(M) can include any token from SEP\_SET(M) or ALPHA\_SET(M), and if the complex NT match is nonempty, then any token starting FIRST(tt ...) could work too. The last piece to consider is  $\varepsilon$ . SEP\_SET(M) and FIRST(tt ...) \ { $\varepsilon$ } cannot contain  $\varepsilon$ , but ALPHA\_SET(M) could. Hence, this definition allows M to accept  $\varepsilon$  if and only if  $\varepsilon \in$  ALPHA\_SET(M) does. This is correct because for M to accept  $\varepsilon$  in the complex NT case, both the complex NT and  $\alpha$  must accept it. If OP = +, meaning that the complex NT cannot be empty, then by definition  $\varepsilon \notin$  ALPHA\_SET(M). Otherwise, the complex NT can accept zero repetitions, and then ALPHA\_SET(M) = FOLLOW( $\alpha$ ). So this definition is correct with respect to \varepsilon as well.

r[macro.ambiguity.sets.def.last]

#### LAST

r[macro.ambiguity.sets.def.last.intro] LAST(M), defined by case analysis on M itself (a sequence of token-trees):

r[macro.ambiguity.sets.def.last.empty]

- if M is the empty sequence, then LAST(M) = { ε } r[macro.ambiguity.sets.def.last.token]
- if M is a singleton token t, then LAST(M) = { t } r[macro.ambiguity.sets.def.last.rep-star]
- if M is the singleton complex NT repeating zero or more times, M = \$( tt ... ) \*, or M = \$( tt ... ) SEP \*
  - Let sep\_set = { SEP } if SEP present; otherwise sep\_set = {}.

- if  $\epsilon \in LAST(tt ...)$  then  $LAST(M) = LAST(tt ...) \cup sep_set$
- otherwise, the sequence tt ... must be non-empty;  $LAST(M) = LAST(tt ...) \cup \{\epsilon\}.$

r[macro.ambiguity.sets.def.last.rep-plus]

- if M is the singleton complex NT repeating one or more times, M =
  \$( tt ... ) +, or M = \$( tt ... ) SEP +
  - Let sep\_set = { SEP } if SEP present; otherwise sep\_set = {}.
  - if  $\varepsilon \in LAST(tt ...)$  then  $LAST(M) = LAST(tt ...) \cup sep_set$
  - otherwise, the sequence tt ... must be non-empty;
     LAST(M) = LAST(tt ...)

r[macro.ambiguity.sets.def.last.rep-question]

 if M is the singleton complex NT repeating zero or one time, M = \$( tt ...) ?, then LAST(M) = LAST(tt ...) U {ε}.

r[macro.ambiguity.sets.def.last.delim]

 if M is a delimited token-tree sequence OPEN tt ... CLOSE, then LAST(M) = { CLOSE }.

r[macro.ambiguity.sets.def.last.sequence]

- if M is a non-empty sequence of token-trees tt uu ...,
  - If  $\varepsilon \in LAST(uu ...)$ , then  $LAST(M) = LAST(tt) \cup (LAST(uu ...) \setminus \{ \varepsilon \})$ .
  - Otherwise, the sequence uu ... must be non-empty; then LAST(M) = LAST(uu ...).

# **Examples of FIRST and LAST**

Below are some examples of FIRST and LAST. (Note in particular how the special  $\epsilon$  element is introduced and eliminated based on the interaction between the pieces of the input.)

Our first example is presented in a tree structure to elaborate on how the analysis of the matcher composes. (Some of the simpler subtrees have been elided.)

INPUT: \$( \$d:ident \$e:expr );\* \$( \$( h )\* );\* \$( f ;)+ g ~~~~~ ~~~~~~ ~ FIRST: { \$d:ident } { \$e:expr } { h } \$( \$d:ident \$e:expr );\* \$(\$( h )\* );\* \$( f INPUT: ; )+ FIRST: { \$d:ident } { h, ε } { f } INPUT: \$( \$d:ident \$e:expr );\* \$( \$( h )\* );\* \$( f ;)+ g ~~~~~~~~~~~ | | FIRST: { $\$d:ident, \varepsilon$ } { h,  $\varepsilon$ , ; } { f } { g } \$d:ident \$e:expr );\* \$( \$( h )\* );\* INPUT: \$( \$( f ; )+ g

FIRST: { \$d:ident, h, ;, f }

Thus:

• FIRST(\$(\$d:ident \$e:expr );\* \$( \$(h)\* );\* \$( f ;)+ g) = {
 \$d:ident, h, ;, f }

~~~~~~~~~~~

Note however that:

FIRST(\$(\$d:ident \$e:expr );\* \$( \$(h)\* );\* \$(\$( f ;)+ g)\*) =
 { \$d:ident, h, ;, f, ε }

Here are similar examples but now for LAST.

- LAST(\$d:ident \$e:expr) = { \$e:expr }
- LAST(\$( \$d:ident \$e:expr );\*) = { \$e:expr, ε}
- LAST(\$( \$d:ident \$e:expr );\* \$(h)\*) = { \$e:expr, ε, h }
- LAST(\$( \$d:ident \$e:expr );\* \$(h)\* \$( f ;)+) = { ; }
- LAST(\$( \$d:ident \$e:expr );\* \$(h)\* \$( f ;)+ g) = { g }

r[macro.ambiguity.sets.def.follow]

# FOLLOW(M)

r[macro.ambiguity.sets.def.follow.intro] Finally, the definition for FOLLOW(M) is built up as follows. pat, expr, etc. represent simple nonterminals with the given fragment specifier.

r[macro.ambiguity.sets.def.follow.pat]

• FOLLOW(pat) = { =>, , , =, |, if, in }`.

r[macro.ambiguity.sets.def.follow.expr-stmt]

FOLLOW(expr) = FOLLOW(expr\_2021) = FOLLOW(stmt) = { =>,
 , ; }`.

r[macro.ambiguity.sets.def.follow.ty-path]

FOLLOW(ty) = FOLLOW(path) = { { , [, , , =>, :, =, >, >>, ; , |, as, where, block nonterminals }.

r[macro.ambiguity.sets.def.follow.vis]

 FOLLOW(vis) = { , l any keyword or identifier except a non-raw priv; any token that can begin a type; ident, ty, and path nonterminals}.

r[macro.ambiguity.sets.def.follow.simple]

• FOLLOW(t) = ANYTOKEN for any other simple token, including block, ident, tt, item, lifetime, literal and meta simple nonterminals, and all terminals.

r[macro.ambiguity.sets.def.follow.other-matcher]

FOLLOW(M), for any other M, is defined as the intersection, as t ranges over (LAST(M) \ {ε}), of FOLLOW(t).

r[macro.ambiguity.sets.def.follow.type-first] The tokens that can begin a type are, as of this writing, {(, [, !, \*, &, &&, ?, lifetimes, >, >>, ::, any non-keyword identifier, super, self, Self, extern, crate, \$crate, \_, for, impl, fn, unsafe, typeof, dyn }, although this list may not be complete because people won't always remember to update the appendix when new ones are added.

Examples of FOLLOW for complex M:

- FOLLOW(\$( \$d:ident \$e:expr )\*) = FOLLOW(\$e:expr)
- FOLLOW(\$( \$d:ident \$e:expr )\* \$(;)\*) = FOLLOW(\$e:expr)
   ∩ ANYTOKEN = FOLLOW(\$e:expr)
- FOLLOW(\$( \$d:ident \$e:expr )\* \$(;)\* \$( f |)+) = ANYTOKEN

## **Examples of valid and invalid matchers**

With the above specification in hand, we can present arguments for why particular matchers are legal and others are not.

- (\$ty:ty < foo ,) : illegal, because FIRST( < foo , ) = { < } ⊈</li>
   FOLLOW(ty)
- (\$ty:ty , foo <) : legal, because FIRST( , foo <) = { , } is ⊆ FOLLOW(ty).</li>
- (\$pa:pat \$pb:pat \$ty:ty ,) : illegal, because FIRST(\$pb:pat \$ty:ty ,) = { \$pb:pat } ⊈ FOLLOW(pat), and also FIRST(\$ty:ty ,) = { \$ty:ty } ⊈ FOLLOW(pat).
- ( \$(\$a:tt \$b:tt)\* ; ) : legal, because FIRST(\$b:tt) = {
   \$b:tt } is ⊆ FOLLOW(tt) = ANYTOKEN, as is FIRST(;) = { ; }.
- ( \$(\$t:tt),\*, \$(t:tt),\*) : legal, (though any attempt to actually use this macro will signal a local ambiguity error during expansion).
- (\$ty:ty \$(; not sep)\* -) : illegal, because FIRST(\$(; not sep)\* -) = { ;, - } is not in FOLLOW(ty).
- (\$(\$ty:ty)-+) : illegal, because separator is not in FOLLOW(ty).
- (\$(\$e:expr)\*) : illegal, because expr NTs are not in FOLLOW(expr NT).

# Influences

Rust is not a particularly original language, with design elements coming from a wide range of sources. Some of these are listed below (including elements that have since been removed):

- SML, OCaml: algebraic data types, pattern matching, type inference, semicolon statement separation
- C++: references, RAII, smart pointers, move semantics, monomorphization, memory model
- ML Kit, Cyclone: region based memory management
- Haskell (GHC): typeclasses, type families
- Newsqueak, Alef, Limbo: channels, concurrency
- Erlang: message passing, thread failure, <del>linked thread failure</del>, <del>lightweight concurrency</del>
- Swift: optional bindings
- Scheme: hygienic macros
- C#: attributes
- Ruby: closure syntax, block syntax
- NIL, Hermes: typestate
- <u>Unicode Annex #31</u>: identifier and pattern syntax

# **Test summary**

The following is a summary of the total tests that are linked to individual rule identifiers within the reference.

{{summary-table}}

# Glossary

#### Abstract syntax tree

An 'abstract syntax tree', or 'AST', is an intermediate representation of the structure of the program when the compiler is compiling it.

#### Alignment

The alignment of a value specifies what addresses values are preferred to start at. Always a power of two. References to a value must be aligned. <u>More</u>.

#### Arity

Arity refers to the number of arguments a function or operator takes. For some examples, f(2, 3) and g(4, 6) have arity 2, while h(8, 2, 6) has arity 3. The ! operator has arity 1.

#### Array

An array, sometimes also called a fixed-size array or an inline array, is a value describing a collection of elements, each selected by an index that can be computed at run time by the program. It occupies a contiguous region of memory.

#### **Associated item**

An associated item is an item that is associated with another item. Associated items are defined in <u>implementations</u> and declared in <u>traits</u>. Only functions, constants, and type aliases can be associated. Contrast to a <u>free item</u>.

#### **Blanket implementation**

Any implementation where a type appears <u>uncovered</u>. impl<T> Foo for T, impl<T> Bar<T> for T, impl<T> Bar<Vec<T>> for T, and impl<T> Bar<T> for Vec<T> are considered blanket impls. However, impl<T> Bar<Vec<T>> for Vec<T> is not a blanket impl, as all instances of T which appear in this impl are covered by Vec.

# Bound

Bounds are constraints on a type or trait. For example, if a bound is placed on the argument a function takes, types passed to that function must abide by that constraint.

#### Combinator

Combinators are higher-order functions that apply only functions and earlier defined combinators to provide a result from its arguments. They can be used to manage control flow in a modular fashion.

## Crate

A crate is the unit of compilation and linking. There are different <u>types</u> <u>of crates</u>, such as libraries or executables. Crates may link and refer to other library crates, called external crates. A crate has a self-contained tree of <u>modules</u>, starting from an unnamed root module called the crate root. <u>Items</u> may be made visible to other crates by marking them as public in the crate root, including through <u>paths</u> of public modules. <u>More</u>.

# Dispatch

Dispatch is the mechanism to determine which specific version of code is actually run when it involves polymorphism. Two major forms of dispatch are static dispatch and dynamic dispatch. While Rust favors static dispatch, it also supports dynamic dispatch through a mechanism called 'trait objects'.

# Dynamically sized type

A dynamically sized type (DST) is a type without a statically known size or alignment.

#### Entity

An *<u>entity</u>* is a language construct that can be referred to in some way within the source program, usually via a <u>path</u>. Entities include <u>types</u>, <u>items</u>,

<u>generic parameters</u>, <u>variable bindings</u>, <u>loop labels</u>, <u>lifetimes</u>, <u>fields</u>, <u>attributes</u>, and <u>lints</u>.

## Expression

An expression is a combination of values, constants, variables, operators and functions that evaluate to a single value, with or without side-effects.

For example, 2 + (3 \* 4) is an expression that returns the value 14.

### Free item

An <u>item</u> that is not a member of an <u>implementation</u>, such as a *free function* or a *free const*. Contrast to an <u>associated item</u>.

# **Fundamental traits**

A fundamental trait is one where adding an impl of it for an existing type is a breaking change. The Fn traits and Sized are fundamental.

### **Fundamental type constructors**

A fundamental type constructor is a type where implementing a <u>blanket</u> <u>implementation</u> over it is a breaking change. &, &mut, Box, and Pin are fundamental.

Any time a type T is considered <u>local</u>, &T, &mut T, Box<T>, and Pin<T> are also considered local. Fundamental type constructors cannot <u>cover</u> other types. Any time the term "covered type" is used, the T in &T, &mut T, Box<T>, and Pin<T> is not considered covered.

# Inhabited

A type is inhabited if it has constructors and therefore can be instantiated. An inhabited type is not "empty" in the sense that there can be values of the type. Opposite of <u>Uninhabited</u>.

#### **Inherent implementation**

An <u>implementation</u> that applies to a nominal type, not to a trait-type pair. <u>More</u>.

## **Inherent method**

A <u>method</u> defined in an <u>inherent implementation</u>, not in a trait implementation.

# Initialized

A variable is initialized if it has been assigned a value and hasn't since been moved from. All other memory locations are assumed to be uninitialized. Only unsafe Rust can create a memory location without initializing it.

# Local trait

A trait which was defined in the current crate. A trait definition is local or not independent of applied type arguments. Given trait Foo<T, U>, Foo is always local, regardless of the types substituted for T and U.

# Local type

A struct, enum, or union which was defined in the current crate. This is not affected by applied type arguments. struct Foo is considered local, but Vec<Foo> is not. LocalType<ForeignType> is local. Type aliases do not affect locality.

## Module

A module is a container for zero or more <u>items</u>. Modules are organized in a tree, starting from an unnamed module at the root called the crate root or the root module. <u>Paths</u> may be used to refer to items from other modules, which may be restricted by <u>visibility rules</u>. <u>More</u>

#### Name

A *name* is an <u>identifier</u> or <u>lifetime or loop label</u> that refers to an <u>entity</u>. A *name binding* is when an entity declaration introduces an identifier or label associated with that entity. <u>Paths</u>, identifiers, and labels are used to refer to an entity.

#### Name resolution

<u>Name resolution</u> is the compile-time process of tying <u>paths</u>, <u>identifiers</u>, and <u>labels</u> to <u>entity</u> declarations.

#### Namespace

A *namespace* is a logical grouping of declared <u>names</u> based on the kind of <u>entity</u> the name refers to. Namespaces allow the occurrence of a name in one namespace to not conflict with the same name in another namespace.

Within a namespace, names are organized in a hierarchy, where each level of the hierarchy has its own collection of named entities.

#### **Nominal types**

Types that can be referred to by a path directly. Specifically <u>enums</u>, <u>structs</u>, <u>unions</u>, and <u>trait object types</u>.

#### **Dyn-compatible traits**

<u>Traits</u> that can be used in <u>trait object types</u> (dyn Trait). Only traits that follow specific <u>rules</u> are *dyn compatible*.

These were formerly known as *object safe* traits.

#### Path

A *path* is a sequence of one or more path segments used to refer to an <u>entity</u> in the current scope or other levels of a <u>namespace</u> hierarchy.

#### Prelude

Prelude, or The Rust Prelude, is a small collection of items - mostly traits - that are imported into every module of every crate. The traits in the prelude are pervasive.

#### Scope

A <u>scope</u> is the region of source text where a named <u>entity</u> may be referenced with that name.

#### Scrutinee

A scrutinee is the expression that is matched on in match expressions and similar pattern matching constructs. For example, in match  $\times$  { A => 1, B => 2 }, the expression  $\times$  is the scrutinee.

#### Size

The size of a value has two definitions.

The first is that it is how much memory must be allocated to store that value.

The second is that it is the offset in bytes between successive elements in an array with that item type.

It is a multiple of the alignment, including zero. The size can change depending on compiler version (as new optimizations are made) and target platform (similar to how usize varies per-platform).

<u>More</u>.

#### Slice

A slice is dynamically-sized view into a contiguous sequence, written as [T].

It is often seen in its borrowed forms, either mutable or shared. The shared slice type is &[T], while the mutable slice type is &mut [T], where T represents the element type.

#### Statement

A statement is the smallest standalone element of a programming language that commands a computer to perform an action.

## **String literal**

A string literal is a string stored directly in the final binary, and so will be valid for the 'static duration.

Its type is 'static duration borrowed string slice, &'static str.

## **String slice**

A string slice is the most primitive string type in Rust, written as str. It is often seen in its borrowed forms, either mutable or shared. The shared string slice type is &str, while the mutable string slice type is &mut str.

Strings slices are always valid UTF-8.

## Trait

A trait is a language item that is used for describing the functionalities a type must provide. It allows a type to make certain promises about its behavior.

Generic functions and generic structs can use traits to constrain, or bound, the types they accept.

### Turbofish

Paths with generic parameters in expressions must prefix the opening brackets with a ::. Combined with the angular brackets for generics, this looks like a fish ::<> . As such, this syntax is colloquially referred to as turbofish syntax.

Examples:

```
let ok_num = 0k::<_, ()>(5);
let vec = [1, 2, 3].iter().map(|n| n * 2).collect::<Vec<_>>();
```

This :: prefix is required to disambiguate generic paths with multiple comparisons in a comma-separate list. See <u>the bastion of the turbofish</u> for an example where not having the prefix would be ambiguous.

## **Uncovered type**

A type which does not appear as an argument to another type. For example, T is uncovered, but the T in Vec<T> is covered. This is only relevant for type arguments.

## **Undefined behavior**

Compile-time or run-time behavior that is not specified. This may result in, but is not limited to: process termination or corruption; improper, incorrect, or unintended computation; or platform-specific results. <u>More</u>.

# Uninhabited

A type is uninhabited if it has no constructors and therefore can never be instantiated. An uninhabited type is "empty" in the sense that there are no values of the type. The canonical example of an uninhabited type is the <u>never type</u> !, or an enum with no variants enum Never { }. Opposite of <u>Inhabited</u>.