



# Ignore Crate lib.rs: Foundation Types

## What This File Does

The lib.rs file is the ignore crate's entry point. Beyond the standard module declarations and re-exports, it defines two fundamental types that pervade the entire crate: the `Error` enum for error handling and the `Match<T>` enum for representing filtering decisions.

These types establish patterns that every other module follows. Understanding them first makes the rest of the crate much clearer.

## Section 1: Crate Overview

The module documentation captures the crate's purpose succinctly: "a fast recursive directory iterator that respects various filters such as globs, file types and `.gitignore` files."

Two usage patterns dominate. The simple `Walk::new("./")` provides an iterator that applies all filtering automatically. The advanced `WalkBuilder` enables fine-grained control over every filtering behavior.

The crate exposes three public modules beyond the walker: `gitignore` for parsing and matching gitignore patterns, `types` for file type definitions, and `overrides` for command-line glob handling. Internal modules (`dir`, `pathutil`, `default_types`) remain private implementation details.

See: Companion Code Section 1

## Section 2: The Error Enum – Variant Overview

The Error enum handles every failure mode in the crate. Eight variants cover the cases:

`Partial` wraps a vector of errors. This handles situations where some operations succeed while others fail – like parsing a gitignore file where some lines are valid and others aren't.

`WithLineNumber`, `WithPath`, and `WithDepth` wrap inner errors with additional context. The wrapping is recursive – you can have an error with a path containing an error with a line number.





## Section 5: I/O Error Handling

I/O errors receive special treatment because they're so common and because callers often need to inspect them.

The `is_io` method checks whether an error is purely an I/O error. Like `is_partial`, it recurses through wrappers. A `WithPath` containing an `Io` is still "an I/O error" for this purpose.

The `io_error` method provides borrowed access to the underlying `std::io::Error` if one exists. This enables checking error kinds without ownership transfer. The companion `into_io_error` consumes `self` for cases where ownership is needed.

The `Clone` implementation for `Error` specially handles the `Io` case. `std::io::Error` isn't `Clone`, but you can reconstruct one from the raw OS error code or by copying the kind and message.

See: Companion Code Section 5

## Section 6: The `PartialErrorBuilder`

The `PartialErrorBuilder` is a private helper for accumulating errors during operations that can partially fail.

Its methods (`push`, `push_ignore_io`, `maybe_push`) provide convenient ways to accumulate errors. The `push_ignore_io` variant silently drops I/O errors – useful when I/O failures should be warnings rather than errors.

The `into_error_option` method collapses the accumulated errors. No errors returns `None`. One error returns that error directly. Multiple errors wraps them in `Partial`.

This builder appears throughout the crate. Any operation that processes multiple items and can fail per-item uses it.

See: Companion Code Section 6

## Section 7: The `Match` Enum – The Filtering Decision

The `Match` enum is how the `ignore` crate represents filtering decisions. Every pattern match, every type check, every override evaluation returns a `Match`.





