



Create a real-time voice incident narrator

Create your own real-time voice assistant that detects incidents by monitoring log files and speaks out loud about issues.

TESTED. GUIDED. WORKING CODE

Contents

About This Book	3
How to Read This Book?	4
Formatting Conventions	5
Introduction	6
Build Log Watcher	7
Initialize Log Watcher	7
Implement Tail Reader	21
Integrate Async Watcher	33
Implement Line Processing	39
Filter And Route Lines	43
Generate Log Summaries	54
Initialize Log Summarizer	54
Run Tasks Concurrently	60
Handle Interrupt Signal	66
Handle Summarizer Events	69
Implement Message Buffer	77
Extend Summarization Buffer	89
Construct Summary Messages	98
Define System Prompt	107
Build Chat Messages	112
Configure Completion Request	121
Configure OpenAI Client	129
Wire AI Summarization Flow	134
Manage Processed Messages	146

About This Book

This section provides essential information about the book's structure, conventions, and how to get the most out of your learning experience.

How to Read This Book?

This book is a **step-by-step guide** to building a real Rust program from scratch. It's designed like an **instruction manual** for a construction kit — if you follow each step carefully, you'll assemble a working solution and gain hands-on development experience.

Step-by-Step Instructions

Each chapter breaks down the development process into clear, manageable steps. Every step shows you exactly what changes to make, which files to modify, and what the result should be. Think of it as building with LEGO blocks — one piece at a time, following the instructions, until you have a complete working program.

Adapted for Learning

Every step includes detailed explanations of all concepts involved. You learn by doing, similar to how AI models are trained — by repetition and practice. Your brain processes each pattern, each solution,

building neural pathways through hands-on experience.

This approach ensures you don't just copy code — you understand **why** it works and **how** the pieces fit together.

Adapted for Vibe-Coding

This book embraces modern **vibe-coding** workflows where you work alongside AI assistants. The premium version includes carefully crafted **prompts** designed for coding agents like Claude, GitHub Copilot, and other AI tools.

These prompts look like this:



*This is a prompt for a coding agent
— copy and use it with AI assistants*

Just copy the prompt, paste it into your AI assistant, and get context-aware code suggestions. This makes development faster while maintaining the learning benefits — you see how AI interprets requirements and produces solutions.

File Changes with Guarantees

At each step, you'll see the **delta of file changes** — exactly what lines were added, modified, or removed. This approach guarantees that if you replicate everything as shown, your program will compile successfully.

No guesswork, no missing pieces — just follow the changes, and you'll have a working program. This makes the book perfect for both manual coding and AI-assisted development.

Who Is This Book For?

Whether you're learning Rust for the first time, exploring AI-assisted coding, or just want a structured project to build, this book provides a clear path from empty project to working software.

Formatting Conventions

This book uses consistent formatting conventions to help you navigate and understand the material more easily.

Code Blocks

Code examples appear in monospace font with syntax highlighting. Line numbers on the left help you reference specific lines:

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

File Changes

When modifying existing files, changed lines are highlighted to show exactly what was added or modified. This helps you track incremental changes throughout the book.

Command Prompts

Terminal commands are shown with a dollar sign prefix and include the current working directory:

```
sh
cargo new my-project
cd my-project
cargo run
```

Emphasis and Formatting

Italic text indicates new terms or concepts being introduced for the first time.

Bold text highlights important library names, framework names, or critical concepts you should pay special attention to.

Inline code appears in monospace font for variable names, function names, file paths, and short code snippets within regular text.

Notes and Callouts

Important information, tips, and warnings appear in special formatted boxes to draw your attention to critical details or common pitfalls.

File Paths

File paths use forward slashes and are shown relative to the project root unless otherwise specified. For example: `src/main.rs` refers to the main source file.

Version Numbers

Code examples in this book are tested with specific versions of Rust and dependencies. While newer versions should generally work, check the book's introduction for the recommended version information.

Introduction

The tool we've built is like an incident radio for your backend: it listens to your logs, notices when things start going wrong, and tells you the story out loud. Instead of staring at scrolling text all day, you get short spoken updates about what's breaking and how it's evolving over time.

Normally, you might do something like:

```
sh
tail -f sample.log | grep -E
'ERROR|WARN'
```

That shows every raw line but doesn't answer the real question: "What is actually happening to my system right now?" With this tool, you can run:

```
sh
cargo run -- sample.log
```

and hear something like: "Login failures are rising, the cache is unstable, and traffic is falling back to the database."

Conceptually, it's a log companion that turns noisy error streams into calm, pe-

riodic status reports. It keeps a rolling memory of past problems for context, focuses on the newest failures, and turns that into a short, non-technical summary you could play over speakers while you work on something else. The idea is that you only drop back into dashboards or log searches when the radio tells you there's a real story unfolding.

We aim this at humans first, not machines, so the summaries are brief, plain-language, and easy to follow even if you're not deeply technical. That also makes it safe to treat as background audio: you can let it run alongside your usual tools and only react when it says something important.

Build Log Watcher

Initialize Log Watcher

Let's begin building a utility to reproduce errors from logs by creating a new binary crate with the `cargo new` command, specifying the `--bin` flag and the crate name.

```
. SH
$ cargo new --bin incident-narrator

Cargo.toml TOML
1 [package]
2 name = "incident-narrator"
3 version = "0.1.0"
4 edition = "2021"
5 [dependencies]

src/main.rs RUST
1 fn main() {
2 }
```

Let's create the first module responsible for reading files and monitoring changes to them. To do this, create a `watcher` module and a struct with the same name inside it.

`src/main.rs`

RUST

```
1 mod watcher;
```

`src/watcher.rs`

RUST

```
1 pub struct Watcher {  
2 }
```


An instance of `Watcher` will monitor exactly one file, but to find and read that file, we need its path.

The most convenient type to store the path inside the struct is `PathBuf`. Import it from the `std::path` module and add a `log_path` field to the struct.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
2 pub struct Watcher {
3     log_path: PathBuf,
4 }
```

Let's create a constructor for the `Watcher` structure: add a `new()` method that takes the path of the file to be watched as its parameter.

You can use the already imported `PathBuf` as the argument type.

`src/watcher.rs`

RUST

```
5 impl Watcher {  
6     pub fn new(log_path: PathBuf) {  
7     }  
8 }
```

Store the provided value in a field of the newly created `Watcher` struct instance, and return it from the `new()` method using `Self` as the return type—an alias for the type itself.

`src/watcher.rs`

RUST

```
5 impl Watcher {  
6     pub fn new(log_path: PathBuf) -> Self {  
7         Self {  
8             log_path,  
9         }  
10    }  
11 }
```

Now we need to get the filename from the command-line parameter and pass it to our [Watcher](#).

To parse the provided arguments more conveniently, import the `clap` crate. Enable the `derive` feature, which lets us turn a struct (that we will add next) into an argument parser.

You can also enable the `env` feature right away so that later we can read the AI model's API key from an environment variable.

Cargo.toml

TOML

```
5 [dependencies]
6 clap = { version = "4.5", features = ["derive", "env"] }
```

Import the `Parser` trait from the `clap` crate, declare a new `Args` struct, and enable parsing by adding a `derive` attribute with the imported type. This creates a command-line argument parser.

The created parser is empty and, for now, it does not accept any arguments.

`src/main.rs`

RUST

```
2 use clap::Parser;
3 #[derive(Parser)]
4 struct Args {
5 }
```

We need the log file name as an argument, which we will watch; to add this field, import the `PathBuf` type we used earlier in the `Watcher` itself, and add the `log_file` field, which is a required positional argument.

In this case, invoke the command by passing the file name.

src/main.rs

RUST

```
3 use std::path::PathBuf;
4 #[derive(Parser)]
5 struct Args {
6     #[arg(value_name = "LOG_FILE")]
7     log_file: PathBuf,
8 }
```

The `Parser` trait that we imported and implemented for the `Args` structure provides the `parse()` function, which parses the command-line arguments and maps them to the structure, or displays help and the required parameters if anything is specified incorrectly. Call this function and store the result in the `args` variable.

Now that we have the log file name, create an instance of the `Watcher` component, import it from the module we created, call the `new()` function we implemented, passing the value of the `log_file` field extracted from the arguments structure as the parameter, and store the resulting instance in the `watcher` variable.

src/main.rs

RUST

```
4 use watcher::Watcher;
10 fn main() {
11     let args = Args::parse();
12     let watcher = Watcher::new(args.log_file);
13 }
```

Now let's add the `run()` method to our `Watcher`, which will serve as the entry point for monitoring file changes and reading the file.

Although the method is still empty, we can already call it on our instance.

src/main.rs

RUST

```
10 fn main() {  
12     let watcher = Watcher::new(args.log_file);  
13     watcher.run();  
14 }
```

src/watcher.rs

RUST

```
5 impl Watcher {  
11     pub fn run(self) {  
12     }  
13 }
```


Even though the program reads log files, it's useful to add logging to the program itself to understand what's happening and when different components start. Later, we will have one task reading and watching the log file for changes, and another sending asynchronous requests to the model to generate summaries.

The `tracing` crate provides macros for logging, and the `tracing-subscriber` crate outputs those logs.

In the latter crate, it's helpful to enable the `env-filter` feature, which lets you control the logging level by setting a filter, for example via the `RUST_LOG` environment variable.

Cargo.toml

TOML

```
5 [dependencies]
6 clap = { version = "4.5", features = ["derive", "env"] }
7 tracing = "0.1"
8 tracing-subscriber = { version = "0.3", features = ["env-
  filter"] }
```

The `tracing-subscriber` crate provides the `fmt()` function, which returns a `SubscriberBuilder` instance that lets you configure logging output.

We don't need to dive into configuration details for this logger and can use the default settings by calling `init()`.

Logs will now appear in the console, so indicate the program start using the `info!` macro from the `tracing` crate.

Also, in the `Watcher` component's `run()` method, use the same macro to report which file path the task is watching.

src/main.rs

RUST

```
10 fn main() {  
12     tracing_subscriber::fmt().init();  
13     tracing::info!("Starting incident narrator...");  
16 }
```

src/watcher.rs

RUST

```
5 impl Watcher {  
11     pub fn run(self) {  
12         tracing::info!("Watching log file: {}",  
self.log_path.display());  
13     }  
14 }
```

It would also be helpful to generate a sample log file to test the program.

Add sample log entries and fictional errors to the file `sample.log`.

You can use the version provided here, or generate a new file by prompting your favorite AI code assistant.

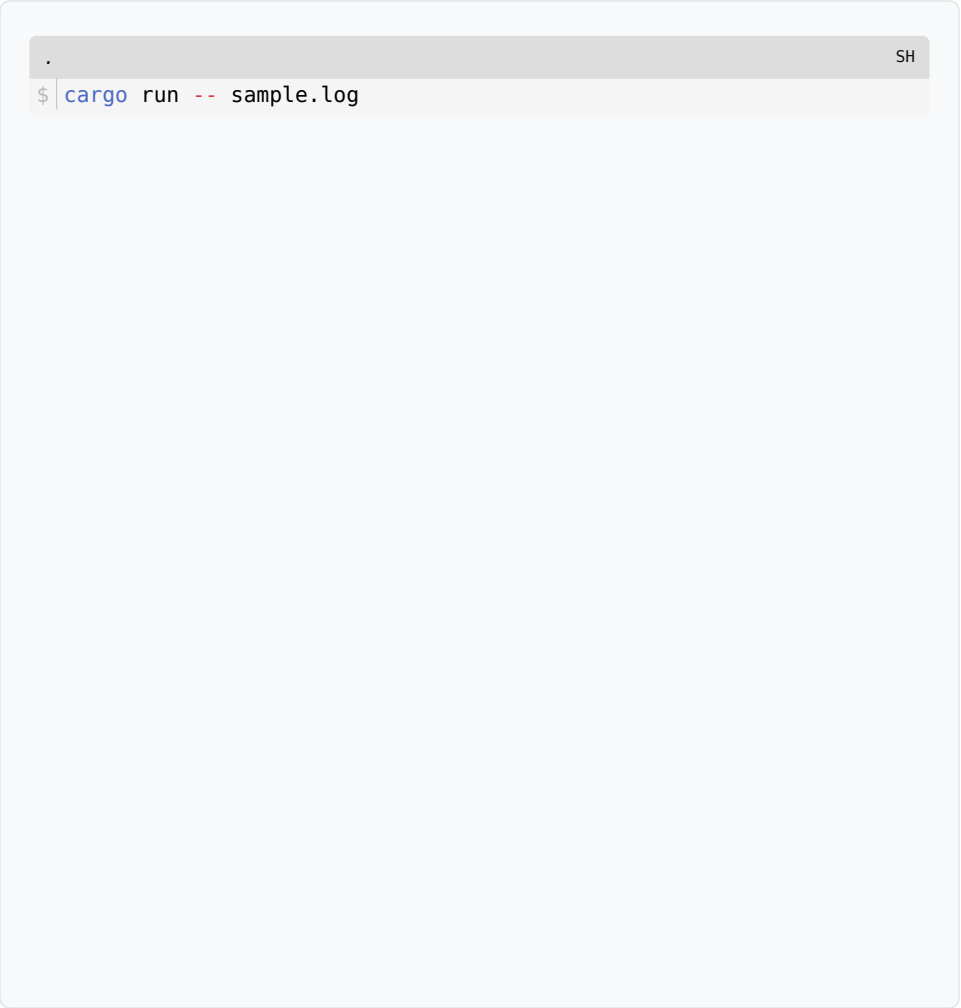
`sample.log`

LOG

```
1 2035-11-16T10:30:15Z INFO Starting application server
2 2035-11-16T10:30:18Z ERROR Database connection timeout
3 2035-11-16T10:30:24Z CRITICAL Cannot start without database
4 2035-11-16T10:30:26Z INFO Connected to backup database
5 2035-11-16T10:30:27Z INFO HTTP server started on port 8080
6 2035-11-16T10:30:32Z ERROR Authentication failed for user 'admin'
7 2035-11-16T10:30:36Z ERROR Failed to process record: missing
  field 'email'
8 2035-11-16T10:30:45Z ERROR Redis connection lost
9 2035-11-16T10:30:46Z WARN Falling back to direct database queries
10 2035-11-16T10:30:47Z FATAL Out of memory
```

Try running the program with `cargo run`, passing the name of our sample log file `sample.log` as an argument, and remember to include the double dash `--` to separate `cargo` options from our program's arguments.

After startup, you will see an informational message indicating that the `Watcher` component has started, along with the name of the file it intends to watch.

A terminal window with a light gray background. The title bar is dark gray with a white dot on the left and the text 'SH' on the right. The terminal content shows a prompt '\$' followed by the command 'cargo run -- sample.log'.

```
.  
$ cargo run -- sample.log
```

Implement Tail Reader

Let's implement the file reading functionality now. We need to regularly read the tail of the file as new log entries are added. In other words, we need to read the new parts of the log file that we have not yet processed.

There are several ways to implement this. One approach is to subscribe to filesystem events and read changes when an event arrives. However, in our case, periodic file reading is sufficient, since we expect entries to be added to the log regularly. Moreover, this approach is often used by utilities because it is simpler to implement and more portable across different systems—there is no need to rely on a system-specific API.

Add a new `read_tail()` method where we will implement the required reading functionality.

src/watcher.rs

RUST

```
5  impl Watcher {  
14      fn read_tail(&mut self) {  
15      }  
16 }
```

For working with the file system, it is convenient to use an asynchronous approach; this keeps the code concise. This functionality is fully provided by the `tokio` crate, so add it to the project.

Also enable the `full` feature, since we will rely not only on reading files with this crate but also on running tasks in its runtime.

You may know that Rust has no built-in runtime—it is as lean as C. What you write is what you get, with no extra magic, so to make an application asynchronous we must explicitly add a runtime that drives async tasks.

`Cargo.toml`

TOML

```
5 [dependencies]
6 clap = { version = "4.5", features = ["derive", "env"] }
7 tokio = { version = "1.48", features = ["full"] }
8 tracing = "0.1"
9 tracing-subscriber = { version = "0.3", features = ["env-
  filter"] }
```

To read a file's tail, we first need to open it. Use the asynchronous `open()` function of the `File` type, which we must import beforehand from the `fs` module provided by the `tokio` crate.

The function takes the file path as its only parameter; for this, we will use the path previously saved in the `log_path` field, so pass a reference to this field as the argument when calling `open`.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
2 use tokio::{
3     fs::File,
4 };
8 impl Watcher {
17     fn read_tail(&mut self) {
18         File::open(&self.log_path);
19     }
20 }
```

As mentioned, the `open()` function is asynchronous, so you need to use the `await` operator to get its result.

You can also do this within an asynchronous block, so add the `async` keyword to our `read_tail()` function to make it asynchronous as well.

src/watcher.rs

RUST

```
8  impl Watcher {  
17     async fn read_tail(&mut self) {  
18         File::open(&self.log_path).await;  
19     }  
20 }
```


Reading a file can fail, as can many other operations the program performs, so we also need a convenient error-handling mechanism. A great solution for this is to use the [anyhow](#) crate.

Cargo.toml

TOML

```
5 [dependencies]
6 anyhow = "1.0"
7 clap = { version = "4.5", features = ["derive", "env"] }
8 tokio = { version = "1.48", features = ["full"] }
9 tracing = "0.1"
10 tracing-subscriber = { version = "0.3", features = ["env-
    filter"] }
```

Import the `Result` type alias from the `anyhow` crate and use it as the return type of the `read_tail()` function.

Also use the `?` operator to handle the result of `open()`, since it returns a file handle wrapped in `Result`; on error, the function will return early.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
2 use anyhow::Result;
9 impl Watcher {
18     async fn read_tail(&mut self) -> Result<()> {
19         File::open(&self.log_path).await?;
20         Ok(())
21     }
22 }
```

Log file entries are generally represented as separate lines. However, we cannot read individual lines directly from a `File` object. We need a special wrapper that stores and accumulates read data in a buffer until an end of line is detected.

This functionality is already provided by the `BufReader` type, so import it from the `io` module of the `tokio` crate.

Store the result of the `open()` function in the `file` variable, and use it to create a `BufReader` instance by calling `new()` and passing the saved `file` object as the parameter. Assign the result to the `reader` variable; we will now use it for operations on the opened file.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::BufReader,
6 };
10 impl Watcher {
19     async fn read_tail(&mut self) -> Result<()> {
20         let file = File::open(&self.log_path).await?;
21         let reader = BufReader::new(file);
22         Ok(())
23     }
24 }
```

To read only the tail of the file instead of the whole file, record the position where you last stopped and continue from there.

Add a `last_position` field of type `usize` to the struct for this purpose, and initialize it to `0` in the associated `new()` function when creating a `Watcher` instance.

We use `usize` because the read operation returns the number of bytes read as a `usize`, so we can add it to a field of the same type without conversion.

src/watcher.rs

RUST

```
7 pub struct Watcher {
8     log_path: PathBuf,
9     last_position: usize,
10 }
11 impl Watcher {
12     pub fn new(log_path: PathBuf) -> Self {
13         Self {
14             log_path,
15             last_position: 0,
16         }
17     }
26 }
```

A file offset is specified not as a simple byte count from the beginning, but via the `SeekFrom` enumeration, which lets you set the offset relative to the start, end, or current position. This enables more flexible navigation within the file. We will always track the position from the very beginning of the file to discard the part already read.

Import the `SeekFrom` enumeration from the `io` module, construct the `Start` variant with the `last_position` value to set the offset from the start of the file, and store it in the `position` variable.

Since we store the position as a `usize`, we need to convert it to `u64`, which is expected as the offset type for the enumeration in use.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::{BufReader, SeekFrom},
6 };
11 impl Watcher {
21     async fn read_tail(&mut self) -> Result<()> {
22         let file = File::open(&self.log_path).await?;
23         let reader = BufReader::new(file);
24         let position = SeekFrom::Start(self.last_position as
u64);
25         Ok(())
26     }
27 }
```

Now we can use the computed offset to move within the `BufReader`.

For this, the `AsyncSeekExt` trait provides the asynchronous `seek()` method to reposition within the file and expects a `SeekFrom` value as its argument.

Because the method is asynchronous, use `await` to run it, and the `?` operator to handle any errors that may occur, for example, if the end of the file is reached while seeking.

Since the `seek()` method will modify the `BufReader`, add `mut` to the declaration of the `reader` variable.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::{AsyncSeekExt, BufReader, SeekFrom},
6 };
11 impl Watcher {
21     async fn read_tail(&mut self) -> Result<()> {
22         let file = File::open(&self.log_path).await?;
23         let mut reader = BufReader::new(file);
24         let position = SeekFrom::Start(self.last_position as
u64);
25         reader.seek(position).await?;
26         Ok(())
27     }
28 }
```

We plan to call the `read_tail()` method in an infinite loop to regularly read updates from the file. To avoid overloading the system, we need to add an interval to give the log file time to grow.

Create an interval constant of type `Duration`, import `Duration` from the `time` module of the `tokio` crate, and set the interval to `500` milliseconds.

The 500-millisecond value might still be too small, and you may want to set it to 2 or 3 seconds.

`src/watcher.rs`

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::{AsyncSeekExt, BufReader, SeekFrom},
6     time::Duration,
7 };
8 const INTERVAL: Duration = Duration::from_millis(500);
```

Then import the `sleep()` function from the same module. Call it at the end of the method, passing the created constant as the argument; since the function is asynchronous, use `await` to pause the method for the specified number of milliseconds.

This could also be done at the end of the main loop, but in this implementation we prefer to wait for the interval only after a successful read, and on failure immediately attempt to read again.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::{AsyncSeekExt, BufReader, SeekFrom},
6     time::{sleep, Duration},
7 };
13 impl Watcher {
23     async fn read_tail(&mut self) -> Result<()> {
24         let file = File::open(&self.log_path).await?;
25         let mut reader = BufReader::new(file);
26         let position = SeekFrom::Start(self.last_position as
u64);
27         reader.seek(position).await?;
28         sleep(INTERVAL).await;
29         Ok(())
30     }
31 }
```


Integrate Async Watcher

At this point, the `read_tail()` method simply opens the file and moves the read position to the start of the unread section. We also need an additional method to process the remaining tail of the file. Add a new `process_lines` method for this task.

Because the read operation is not guaranteed to succeed, you can use `Result` as the return type to handle all errors conveniently later.

src/watcher.rs

RUST

```
13 impl Watcher {  
31     async fn process_lines(&mut self) -> Result<()> {  
32         Ok(())  
33     }  
34 }
```

Since we need to read the tail of the file, we can keep using the `BufReader` we have advanced to the required position, so just pass a mutable reference to it as a parameter to the new method.

Call this method from `read_tail()` right after calling `seek()`, but before we start waiting with the `sleep()` function. And remember that the new method is asynchronous and may return an error.

src/watcher.rs

RUST

```
13 impl Watcher {
23     async fn read_tail(&mut self) -> Result<()> {
24         let file = File::open(&self.log_path).await?;
25         let mut reader = BufReader::new(file);
26         let position = SeekFrom::Start(self.last_position as
u64);
27         reader.seek(position).await?;
28         self.process_lines(&mut reader).await?;
29         sleep(INTERVAL).await;
30         Ok(())
31     }
32     async fn process_lines(&mut self, reader: &mut
BufReader<File>) -> Result<()> {
33         Ok(())
34     }
35 }
```

Earlier we discussed creating a general infinite file-reading loop that calls the `read_tail()` method. Add a `loop` to the `run` method and call the method from it.

Since the method will update the offset value, mark `self` as `mut`.

```
src/watcher.rs RUST
13 impl Watcher {
20     pub fn run(mut self) {
21         tracing::info!("Watching log file: {}",
self.log_path.display());
22         loop {
23             self.read_tail();
24         }
25     }
38 }
```

The method invoked in the `loop` returns a `Result` and is asynchronous, which directly affects the `run()` method as well. Therefore, make it asynchronous too and return a `Result` to meet the same criteria.

```
src/watcher.rs RUST
13 impl Watcher {
20     pub async fn run(mut self) -> Result<()> {
21         tracing::info!("Watching log file: {}",
22             self.log_path.display());
23         loop {
24             self.read_tail().await?;
25         }
26     }
27 }
```

But that's not all. Although we called the `run()` method, it is now asynchronous. Therefore, we also need to make the entire `main()` function asynchronous.

The easiest way to do this is to use the `main` macro from the `tokio` crate, which will start the runtime and call our asynchronous version of the function.

Just remember to add the `async` keyword to its declaration and use `await` to run the asynchronous `run()` method.

src/main.rs

RUST

```
10 #[tokio::main]
11 async fn main() {
15     let watcher = Watcher::new(args.log_file);
16     watcher.run().await;
17 }
```

You also need to handle any error that may occur while running `Watcher`. In this implementation, we do not plan to retry reading the file if an error occurs during reading, so simply aborting execution is sufficient.

You can do this by handling the result returned by the `run()` method. Use the `?` operator, first importing the `Result` alias from the `anyhow` crate and returning it from the `main()` function. At the end of the function, return `Ok(())` if the read task completes successfully.

As you remember, our `Watcher` has an infinite loop, and the program will not actually exit until it is terminated entirely, but we will improve this a bit later.

src/main.rs

RUST

```
2 use anyhow::Result;
11 #[tokio::main]
12 async fn main() -> Result<()> {
16     let watcher = Watcher::new(args.log_file);
17     watcher.run().await?;
18     Ok(())
19 }
```

Implement Line Processing

Let's return to implementing the `process_lines()` method, which is responsible for reading the tail of the log file.

The first step is to read the next line from the provided `BufReader`. As you recall, we wrapped the file in a `BufReader` precisely because it provides the `read_line()` method for line-by-line reading.

In fact, this method is provided by the `AsyncBufReadExt` trait, which should be imported from the `io` module of the `tokio` crate. This trait is implemented by the `BufReader` type, whose instance we stored in the `reader` variable.

The method takes a mutable reference to a buffer as its argument, where the line will be written. Create a new one of type `String` by calling its `new()` method; this produces an empty string into which `read_line()` will read the next line from the file.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
3 use tokio::{
4     fs::File,
5     io::{AsyncBufReadExt, AsyncSeekExt, BufReader, SeekFrom},
6     time::{sleep, Duration},
7 };
13 impl Watcher {
35     async fn process_lines(&mut self, reader: &mut
        BufReader<File>) -> Result<()> {
36         let mut line = String::new();
37         reader.read_line(&mut line).await?;
38         Ok(())
39     }
40 }
```

We read one line once; to read all the remaining lines, we need to repeat the same operation multiple times.

To do this, move the line reading (the `read_line()` call) into a `loop`. To avoid creating a buffer on each iteration of the loop, reuse the existing one by clearing it to an empty string with the `clear()` method.

```
src/watcher.rs RUST
13 impl Watcher {
35     async fn process_lines(&mut self, reader: &mut
        BufReader<File>) -> Result<()> {
36         let mut line = String::new();
37         loop {
38             line.clear();
39             reader.read_line(&mut line).await?;
40         }
41         Ok(())
42     }
43 }
```


The loop will run indefinitely until a read error occurs; the compiler will even warn that the `Ok()` result at the end of the `process_lines()` method will never be reached.

However, the `read_line()` method returns the number of bytes read, and we can stop further reading when no new data is read (it returns `0`).

Store the number of bytes read in the `bytes_read` variable, and if the value is `0`, break the loop using the `break` statement.

src/watcher.rs

RUST

```
13 impl Watcher {
35     async fn process_lines(&mut self, reader: &mut
BufReader<File>) -> Result<()> {
36         let mut line = String::new();
37         loop {
38             line.clear();
39             let bytes_read = reader.read_line(&mut line).await?;
40             if bytes_read == 0 {
41                 break;
42             }
43         }
44         Ok(())
45     }
46 }
```

But, as you recall, we also store the number of bytes read from the file in the `last_position` field.

It is time to update it—after reading a line, add the number of bytes read to this field, so that on the next iteration we start reading from the next line, since the read cursor will be moved using `seek()` before this method is called.

src/watcher.rs

RUST

```
13 impl Watcher {
35     async fn process_lines(&mut self, reader: &mut
BufReader<File>) -> Result<()> {
36         let mut line = String::new();
37         loop {
38             line.clear();
39             let bytes_read = reader.read_line(&mut line).await?;
40             if bytes_read == 0 {
41                 break;
42             }
43             self.last_position += bytes_read;
44         }
45         Ok(())
46     }
47 }
```

Filter And Route Lines

We have read a line, and it is now stored in a buffer. What should we do with this line next? Obviously, the first step is to filter the lines: separate entries with errors and warnings, and discard the rest.

We can do this by comparing the line to a pattern; the simplest yet most flexible approach is to use *regular expressions*. This functionality is not included in the standard library, but is fully implemented by the [regex](#) crate—add it to the project.

Cargo.toml

TOML

```
5 [dependencies]
6 anyhow = "1.0"
7 clap = { version = "4.5", features = ["derive", "env"] }
8 regex = "1.12"
9 tokio = { version = "1.48", features = ["full"] }
10 tracing = "0.1"
11 tracing-subscriber = { version = "0.3", features = ["env-
    filter"] }
```

Let's define a regular expression pattern that can detect lines with errors and warnings. For now, it can just be a string literal. Below is a detailed description of each part of the regular expression that detects any of the listed words:

Part	Description
<code>r"..."</code>	A raw string in Rust. This means backslashes (<code>\</code>) are not interpreted by Rust as escape sequences but are passed directly to the regular expression.
<code>(...)</code>	Capturing group. This parenthesized group combines the listed words.
<code>error</code>	Searches for the exact sequence of characters "error".
<code>warning</code>	Searches for "warning".
<code>warn</code>	Searches for "warn".
<code>fatal</code>	Searches for "fatal".
<code>critical</code>	Searches for "critical".
<code> </code>	Alternation (OR).

`src/watcher.rs`

RUST

```
8 const INTERVAL: Duration = Duration::from_millis(500);
9 const PATTERN: &str = r"(error|warning|warn|fatal|critical)";
```

To make the pattern more versatile, add case-insensitive matching and mark word boundaries by including the following elements in the pattern:

Part	Description
<code>(?i)</code>	A flag enabling case-insensitive matching. This means the expression will match <code>error</code> , <code>ERROR</code> , <code>Error</code> , <code>eRrOr</code> , etc.
<code>\b</code>	Word boundary. A zero-width assertion that matches the position between a word character (w, i.e., a letter, digit, or underscore) and a non-word character (W), or the start/end of the string. This ensures the match is a whole word, not part of another word (for example, it finds <code>warn</code> but not <code>be-ware</code> or <code>warningly</code>).
<code>\b</code>	Word boundary. Closes the pattern by ensuring that a non-word character or the end of the string follows the matched word.

`src/watcher.rs`

RUST

```
8 const INTERVAL: Duration = Duration::from_millis(500);
9 const PATTERN: &str = r"(?i)\b(error|warning|warn|fatal|
critical)\b";
```

Now we can construct a `Regex` instance from the generated pattern.

Import the `Regex` type from the `regex` crate and use the `new()` method, which takes the pattern as its only parameter; pass it the name of the constant we declared earlier. Store the result in the `regex` variable.

```
src/watcher.rs RUST
1 use std::path::PathBuf;
3 use regex::Regex;
15 impl Watcher {
16     pub fn new(log_path: PathBuf) -> Self {
17         let regex = Regex::new(PATTERN);
18         Self {
19             log_path,
20             last_position: 0,
21         }
22     }
50 }
```

The `regex` variable indeed stores a value of type `Result`. Since we defined the pattern in code, it is enough to verify once that it compiles and then use `unwrap()` to discard the `Result` wrapper.

However, it is still good practice to check the result, since you might change the pattern from time to time, or it may even be provided by the user.

To do this, wrap the return value of the `new()` constructor in a `Result`, wrap the constructed `Self` instance in the successful `Ok` variant, and handle the result of constructing the regular expression with the `?` operator.

`src/main.rs`

RUST

```
11 #[tokio::main]
12 async fn main() -> Result<()> {
16     let watcher = Watcher::new(args.log_file)?;
17     watcher.run().await?;
18     Ok(())
19 }
```

`src/watcher.rs`

RUST

```
15 impl Watcher {
16     pub fn new(log_path: PathBuf) -> Result<Self> {
17         let regex = Regex::new(PATTERN);
18         Ok(Self {
19             log_path,
20             last_position: 0,
21         })
22     }
50 }
```

Move the constructed regular expression into a field of the `Watcher` struct.

To do this, add a new `regex` field of type `Regex`, and initialize it in the `new()` function by setting its value to the constructed regular expression stored in the `regex` variable.

```
src/watcher.rs RUST
11 pub struct Watcher {
12     log_path: PathBuf,
13     last_position: usize,
14     regex: Regex,
15 }
16 impl Watcher {
17     pub fn new(log_path: PathBuf) -> Result<Self> {
18         let regex = Regex::new(PATTERN)?;
19         Ok(Self {
20             log_path,
21             last_position: 0,
22             regex,
23         })
24     }
25 }
```


Use the constructed regular expression in the `process_line()` method to determine whether the read line matches it.

For this purpose, the `Regex` type provides the `is_match()` method, which takes a string parameter to check whether it contains a match for the regular expression. Add an `if` block that checks this condition.

```
src/watcher.rs RUST
16 impl Watcher {
40     async fn process_lines(&mut self, reader: &mut
BufReader<File>) -> Result<()> {
41         let mut line = String::new();
42         loop {
43             line.clear();
44             let bytes_read = reader.read_line(&mut line).await?;
45             if bytes_read == 0 {
46                 break;
47             }
48             self.last_position += bytes_read;
49             if self.regex.is_match(&line) {
50                 }
51         }
52         Ok(())
53     }
54 }
```

We have built a mechanism that detects lines with errors or warnings. Now we need to pass them to a component that can process them into a summary.

This should be a separate entity because we don't want to interrupt reading the update file while a request to the AI model is in progress.

For this purpose, the `mpsc` channel from the `tokio` crate is a perfect fit. Import it from the `sync` module.

And add an `UnboundedSender` as a parameter to the `new()` function, which will be used to send log lines.

src/watcher.rs

RUST

```
1 use std::path::PathBuf;
4 use tokio::{
5     fs::File,
6     io::{AsyncBufReadExt, AsyncSeekExt, BufReader, SeekFrom},
7     sync::mpsc,
8     time::{sleep, Duration},
9 };
17 impl Watcher {
18     pub fn new(log_path: PathBuf, tx:
mpsc::UnboundedSender<String>) -> Result<Self> {
19         let regex = Regex::new(PATTERN)?;
20         Ok(Self {
21             log_path,
22             last_position: 0,
23             regex,
24         })
25     }
55 }
```

Let's create a channel instance right away for `Watcher` to send the lines it reads.

Import the `mpsc` module and use its `unbounded_channel()` function, which returns a tuple with `UnboundedSender` and `UnboundedReceiver`.

Destructure the tuple into `log_tx` and `log_rx`, and pass the `UnboundedSender` stored in `log_tx` as the additional parameter to the `new()` function we just added.

```
src/main.rs RUST
5 use tokio::sync::mpsc;
12 #[tokio::main]
13 async fn main() -> Result<()> {
17     let (log_tx, log_rx) = mpsc::unbounded_channel();
18     let watcher = Watcher::new(args.log_file, log_tx)?;
19     watcher.run().await?;
20     Ok(())
21 }
```

Store the passed `UnboundedSender` in the `Watcher` struct by adding a `tx` field of the appropriate type and initializing it when creating the instance inside the `new()` function.

```
src/watcher.rs RUST
12 pub struct Watcher {
13     log_path: PathBuf,
14     last_position: usize,
15     regex: Regex,
16     tx: mpsc::UnboundedSender<String>,
17 }
18 impl Watcher {
19     pub fn new(log_path: PathBuf, tx:
mpsc::UnboundedSender<String>) -> Result<Self> {
20         let regex = Regex::new(PATTERN)?;
21         Ok(Self {
22             log_path,
23             last_position: 0,
24             regex,
25             tx,
26         })
27     }
28 }
57 }
```

We can now finish implementing the `process_lines()` method by using the channel sender stored in the struct's `tx` field.

If the read log entry matches our pattern (`if` block), call the sender's `send()` method, passing it the cloned string.

We clone the string because the `String` in `line` is reused as a buffer, which is efficient since not all lines need to be sent; therefore, we clone only the lines forwarded for processing, and we do not need to create a new string for reading (its capacity will grow automatically, and only when the next line is larger than the already allocated memory that could hold any of the previous lines).

src/watcher.rs

RUST

```

18 impl Watcher {
43     async fn process_lines(&mut self, reader: &mut
BufReader<File>) -> Result<()> {
44         let mut line = String::new();
45         loop {
46             line.clear();
47             let bytes_read = reader.read_line(&mut line).await?;
48             if bytes_read == 0 {
49                 break;
50             }
51             self.last_position += bytes_read;
52             if self.regex.is_match(&line) {
53                 self.tx.send(line.clone())?;
54             }
55         }
56         Ok(())
57     }
58 }

```

Generate Log Summaries

Initialize Log Summarizer

We will now create an asynchronous task that processes lines with errors and warnings, turning them into a summary.

Add the `summarizer` module to `main.rs` and create the corresponding empty module file.

Then add the `Summarizer` struct to this new file.

```
src/main.rs
```

```
RUST
```

```
1 mod summarizer;  
2 mod watcher;
```

```
src/summarizer.rs
```

```
RUST
```

```
1 pub struct Summarizer {  
2 }
```

Since we previously created a channel for sending messages for processing, add an `UnboundedReceiver` to our new structure to receive messages from this channel. Save it as the structure's `rx` field.

And of course, import the `mpsc` module from the `tokio` crate's `sync` module, which contains the type we need.

`src/summarizer.rs`

RUST

```
1 use tokio::{
2     sync::mpsc,
3 };
4 pub struct Summarizer {
5     rx: mpsc::UnboundedReceiver<String>,
6 }
```

Add an implementation block for the `Summarizer` struct and declare the associated `new()` function in it to create an instance of the struct.

```
src/summarizer.rs RUST  
7 impl Summarizer {  
8     pub fn new(  
9     ) {  
10    }  
11 }
```


The method should take an `UnboundedReceiver` as a parameter and return an instance of the created structure; we can use the `Self` alias for this.

And in the function body, we should create a `Summarizer` instance, filling its single `rx` field with the passed value.

src/summarizer.rs

RUST

```
7 impl Summarizer {  
8     pub fn new(  
9         rx: mpsc::UnboundedReceiver<String>,  
10    ) -> Self {  
11        Self {  
12            rx,  
13        }  
14    }  
15 }
```

Now add an asynchronous `run()` method to the struct that will pull messages from the channel and process them.

It is also good practice to return a `Result` from the `anyhow` crate and return `Ok` on successful completion of the method. Using `Result` makes error handling convenient and is useful for entry points of asynchronous tasks.

src/summarizer.rs

RUST

```
1 use anyhow::Result;
8 impl Summarizer {
16     pub async fn run(mut self) -> Result<()> {
17         Ok(())
18     }
19 }
```

As you recall, we used the `tracing` crate to indicate the startup of the `Watcher` component. We can do the same for the new `Summarizer`.

Use the `info!` macro to report the start of the execution of the `run()` method. This way you will see that the message processing task has started successfully.

src/summarizer.rs

RUST

```
8  impl Summarizer {  
16     pub async fn run(mut self) -> Result<()> {  
17         tracing::info!("Starting summarizer");  
18         Ok(())  
19     }  
20 }
```

Run Tasks Concurrently

Since we now have two asynchronous tasks, [Watcher](#) and [Summarizer](#), we need to run them together. The [futures](#) crate can help with this, so add it to your dependencies.

Cargo.toml

TOML

```
5 [dependencies]
6 anyhow = "1.0"
7 clap = { version = "4.5", features = ["derive", "env"] }
8 futures = "0.3.31"
9 regex = "1.12"
10 tokio = { version = "1.48", features = ["full"] }
11 tracing = "0.1"
12 tracing-subscriber = { version = "0.3", features = ["env-
  filter"] }
```

Because Rust is strongly typed, to run multiple different tasks we need to bring them to a single type.

For an asynchronous task, we can use the `boxed()` method of the `FutureExt` trait from the `future` module of the `futures` crate. This method boxes the asynchronous task into a `Box` (i.e., allocates it on the heap), producing a value that implements the `Future` trait.

src/main.rs

RUST

```
5 use futures::future::FutureExt;
14 #[tokio::main]
15 async fn main() -> Result<()> {
20     let watcher = Watcher::new(args.log_file, log_tx)?;
21     watcher.run().boxed().await?;
22     Ok(())
23 }
```

To run a group of tasks concurrently, we need to collect them into a `Vec`. Create a new mutable vector `tasks` and add the boxed asynchronous task `Watcher` to this list.

By the way, you no longer need to wait for the task to complete with `await`, since we will do that next for the entire group of tasks at once.

src/main.rs

RUST

```
14 #[tokio::main]
15 async fn main() -> Result<()> {
20     let mut tasks = Vec::new();
21     let watcher = Watcher::new(args.log_file, log_tx)?;
22     tasks.push(watcher.run().boxed());
23     Ok(())
24 }
```

Even though there is only one task in the list for now, let's execute the task list right away.

Import the `select_all()` function from the `future` module. Call it and pass our `tasks` list as the parameter. And since the function is asynchronous, apply `await` to execute it.

The `select_all()` function returns a tuple whose first element (index `0`) is the result of the earliest completed asynchronous task, and whose second element is all unfinished tasks that can continue running.

Since all our asynchronous tasks return a `Result`, we can immediately access index `0`—the result of the first completed task—and use the `?` operator to handle an error if the asynchronous task completed with one.

src/main.rs

RUST

```
5 use futures::future::{select_all, FutureExt};
14 #[tokio::main]
15 async fn main() -> Result<()> {
20     let mut tasks = Vec::new();
23     select_all(tasks).await.0?;
24     Ok(())
25 }
```

Now let's create an instance of the `Summarizer` task. Import this struct from the `summarizer` module and use the associated `new()` function, passing the `UnboundedReceiver` stored in the `log_rx` variable as a parameter so it can receive log file entries.

```
src/main.rs RUST  
7 use summarizer::Summarizer;  
15 #[tokio::main]  
16 async fn main() -> Result<()> {  
21     let mut tasks = Vec::new();  
24     let summarizer = Summarizer::new(log_rx);  
25     select_all(tasks).await.0?;  
26     Ok(())  
27 }
```


Call the `run()` method on the created structure to turn it into an asynchronous task. Wrap it in a `Box` using the `boxed()` method, and add it to the `tasks` list by calling its `push()` method.

As a result, the list contains two tasks, and both will run concurrently.

```
src/main.rs RUST  
15 #[tokio::main]  
16 async fn main() -> Result<()> {  
21     let mut tasks = Vec::new();  
24     let summarizer = Summarizer::new(log_rx);  
25     tasks.push(summarizer.run().boxed());  
26     select_all(tasks).await.0?;  
27     Ok(())  
28 }
```

Handle Interrupt Signal

Let's also explicitly intercept the `SIGINT` interrupt signal, which can be sent from the terminal to the process with `Ctrl+C`. This will allow us to gracefully finish the `select_all()` call and shut down the program.

To do this, import the `ctrl_c()` function from the `signal` module of the `tokio` crate. Then call it to create an async future that listens for the interrupt signal, and store it in the `interrupt` variable.

src/main.rs

RUST

```
8 use tokio::{signal::ctrl_c, sync::mpsc};
15 #[tokio::main]
16 async fn main() -> Result<()> {
21     let mut tasks = Vec::new();
26     let interrupt = ctrl_c();
27     select_all(tasks).await.0?;
28     Ok(())
29 }
```

Add the created task to the `tasks` list by first calling its `boxed()` method to convert it to the common type stored in the list: a `Future` inside a `Box`.

However, this *code will not compile* because the output type of this `Future` still differs from the others in our task list.

```
src/main.rs RUST
15 #[tokio::main]
16 async fn main() -> Result<()> {
21     let mut tasks = Vec::new();
26     let interrupt = ctrl_c();
27     tasks.push(interrupt.boxed());
28     select_all(tasks).await.0?;
29     Ok(())
30 }
```

The mismatch is that our code returns `Error` from the `anyhow` crate, while `ctrl_c()` returns `std::io::Error`.

We can resolve this by explicitly importing `Error` from `anyhow` and using the `map_err()` method provided by the `TryFutureExt` trait, which should be imported from the `future` module.

This trait is a convenient layer for working with `Future` objects that return `Result` (specified via the `Future` trait's associated `Output` type).

The `map_err()` method takes a function that converts one error type to another. Therefore, we can directly use the `from()` function on the `Error` type from the `anyhow` crate.

Since the returned types are now fully compatible, the code compiles successfully.

src/main.rs

RUST

```
3 use anyhow::{Error, Result};
5 use futures::future::{select_all, FutureExt, TryFutureExt};
15 #[tokio::main]
16 async fn main() -> Result<> {
21     let mut tasks = Vec::new();
26     let interrupt = ctrl_c().map_err(Error::from);
27     tasks.push(interrupt.boxed());
28     select_all(tasks).await.0?;
29     Ok(())
30 }
```

Handle Summarizer Events

Let's return to the `Summarizer` implementation and add an asynchronous `step()` method to it.

In that method, we will process one incoming event at a time and later call it in a loop inside the `run()` method.

`src/summarizer.rs`

RUST

```
8  impl Summarizer {  
20      async fn step(&mut self) {  
21          }  
22  }
```

In the `step()` method, we need to receive a message from an `UnboundedReceiver` instance; you can do this with the `recv()` method.

The method is asynchronous, so if we wait for it with `await`, we lose the ability to react to other events.

A better approach here is to use the `select!` macro, which allows us to handle multiple event streams. Import this macro from the `tokio` crate.

Use the macro to add a branch that reads messages from the receiver, and store the read message in the `message` variable.

src/summarizer.rs

RUST

```
2 use tokio::{
3     select,
4     sync::mpsc,
5 };
9 impl Summarizer {
21     async fn step(&mut self) {
22         select! {
23             message = self.rx.recv() => {
24             }
25         }
26     }
27 }
```

You can process incoming messages directly in the `step()` method, but a cleaner approach is to define a dedicated method for each event type.

Therefore, add a `handle_message()` method that processes log messages forwarded by the `Watcher` worker on the exchange channel.

This method should take the message as a parameter (type `String`, the type we send through the channel).

However, because the channel may close, `recv()` returns an `Option` where `None` signals the end of transmission.

We must account for this to interrupt processing, so the parameter should be an `Option<String>`.

src/summarizer.rs

RUST

```
9  impl Summarizer {  
27      fn handle_message(&mut self, message: Option<String>) {  
28      }  
29 }
```

Call the new `handle_message()` method in the `select!` branch, passing the message received from the channel as its argument.

Also use `await` when calling `step()` because it is asynchronous and uses `select!` internally.

```
src/summarizer.rs RUST
9  impl Summarizer {
17      pub async fn run(mut self) -> Result<()> {
18          tracing::info!("Starting summarizer");
19          self.step().await;
20          Ok(())
21      }
22      async fn step(&mut self) {
23          select! {
24              message = self.rx.recv() => {
25                  self.handle_message(message);
26              }
27          }
28      }
31 }
```


Let's now add a flag that will control the lifetime of the entire `Summarizer`.

Add an `active` flag of type `bool` to the struct, and initialize it to `true` in the `new()` function.

src/summarizer.rs

RUST

```
6 pub struct Summarizer {
7     rx: mpsc::UnboundedReceiver<String>,
8     active: bool,
9 }
10 impl Summarizer {
11     pub fn new(
12         rx: mpsc::UnboundedReceiver<String>,
13     ) -> Self {
14         Self {
15             rx,
16             active: true,
17         }
18     }
19 }
33 }
```

Now add a `while` loop to the `run()` method that calls `step()` in its body.

The loop's exit condition will be checking the `active` flag, which is `true` at startup.

```
src/summarizer.rs RUST
10 impl Summarizer {
19     pub async fn run(mut self) -> Result<()> {
20         tracing::info!("Starting summarizer");
21         while self.active {
22             self.step().await;
23         }
24         Ok(())
25     }
35 }
```

In the `handle_message()` method implementation, we first check whether a message actually arrived—that is, the `Option` variant is `Some`—and then extract it for processing using the `if let` construct.

```
src/summarizer.rs RUST  
10 impl Summarizer {  
33     fn handle_message(&mut self, message: Option<String>) {  
34         if let Some(msg) = message {  
35             }  
36     }  
37 }
```

Conversely, when the incoming message is `None`, we set the `active` flag to `false`. This will stop the main loop in the `run()` method.

That is, after we receive the last message from the log, the next loop iteration will check this flag and terminate execution.

src/summarizer.rs

RUST

```
10 impl Summarizer {
33     fn handle_message(&mut self, message: Option<String>) {
34         if let Some(msg) = message {
35         } else {
36             self.active = false;
37         }
38     }
39 }
```

Implement Message Buffer

Now we need to accumulate messages coming from logs to build a summary not for a single message—since that makes no sense—but for an entire group.

At the same time, we need to implement a buffer that stores two categories of messages: those we have never processed and those previously processed.

We will use processed messages to remind the model of the overall context of all issues. And new messages will let us report specifically on new problems in the logs of the monitored application.

For this purpose, let's create a separate `buffer` module and declare the `MessageBuffer` struct in it.

`src/buffer.rs`

RUST

```
1 pub struct MessageBuffer {  
2 }
```

`src/main.rs`

RUST

```
1 mod buffer;  
2 mod summarizer;  
3 mod watcher;
```

Add a `new()` function to the implementation of the `MessageBuffer` struct.

The `new()` method will return instances of the struct, and you can create it by constructing an empty struct using the `Self` alias.

```
src/buffer.rs RUST  
3 impl MessageBuffer {  
4     pub fn new() -> Self {  
5         Self {  
6             }  
7     }  
8 }
```

Although `MessageBuffer` does not yet provide useful functionality, we can already add it to our `Summarizer` so we can later decide which functionality (methods) we actually need.

Bring the `MessageBuffer` type into scope from the crate's `buffer` module, and add a `buffer` field of this type to the struct.

Also, in the `new()` function when creating the `Summarizer` struct, initialize this field by calling the `MessageBuffer` struct's `new()` method with no arguments.

src/summarizer.rs

RUST

```
1 use crate::buffer::MessageBuffer;
7 pub struct Summarizer {
8     rx: mpsc::UnboundedReceiver<String>,
9     active: bool,
10    buffer: MessageBuffer,
11 }
12 impl Summarizer {
13     pub fn new(
14         rx: mpsc::UnboundedReceiver<String>,
15     ) -> Self {
16         Self {
17             rx,
18             active: true,
19             buffer: MessageBuffer::new(),
20         }
21     }
22 }
```

We will definitely need a `MessageBuffer` method to add a new message to the buffer.

Create a method `add_new()` for this purpose that takes a `String` message as its parameter.

`src/buffer.rs`

RUST

```
3 impl MessageBuffer {  
8     pub fn add_new(&mut self, message: String) {  
9     }  
10 }
```


An incoming message needs to be stored somewhere. For this purpose, add a `new_messages` field to the buffer struct, of type `Vec` containing `String` objects.

Initialize this field with an empty vector by calling its `new()` method in the `MessageBuffer` constructor.

And immediately add the incoming message in the `add_new()` method to the `new_messages` vector, using its `push()` method to append the message to the end of the list.

src/buffer.rs

RUST

```
1 pub struct MessageBuffer {
2     new_messages: Vec<String>,
3 }
4 impl MessageBuffer {
5     pub fn new() -> Self {
6         Self {
7             new_messages: Vec::new(),
8         }
9     }
10    pub fn add_new(&mut self, message: String) {
11        self.new_messages.push(message);
12    }
13 }
```

Immediately use the buffer's new `add_new()` method to add the incoming message within the `handle_message()` method.

Note that we did not limit the size of the new-messages buffer, so it will grow indefinitely, accumulating more and more messages.

To avoid overflow, we need to periodically process messages from this buffer.

src/summarizer.rs

RUST

```
12 impl Summarizer {  
36     fn handle_message(&mut self, message: Option<String>) {  
37         if let Some(msg) = message {  
38             self.buffer.add_new(msg);  
39         } else {  
40             self.active = false;  
41         }  
42     }  
43 }
```

Let's add a `Duration` constant `INTERVAL` that we'll use as the interval for sending the buffer.

Import the `Duration` type from the `time` module of the `tokio` crate.

Use the `from_secs()` function to set the constant's value. The function takes a number of seconds, so pass `5` as the argument; we will process records every five seconds.

`src/summarizer.rs`

RUST

```
3 use tokio::{
4     select,
5     sync::mpsc,
6     time::Duration,
7 };
8 const INTERVAL: Duration = Duration::from_secs(5);
```

Import from the `time` module the `Interval` type, a struct that measures time intervals and asynchronously waits for them to complete. Also import the `interval()` function that creates an instance of this type.

Add an `interval` field of type `Interval` to the struct, and initialize it by calling `interval()` with a `Duration` parameter—the interval length—using the constant we defined earlier as the value.

`src/summarizer.rs`

RUST

```
3 use tokio::{
4     select,
5     sync::mpsc,
6     time::{interval, Duration, Interval},
7 };
9 pub struct Summarizer {
10     rx: mpsc::UnboundedReceiver<String>,
11     active: bool,
12     buffer: MessageBuffer,
13     interval: Interval,
14 }
15 impl Summarizer {
16     pub fn new(
17         rx: mpsc::UnboundedReceiver<String>,
18     ) -> Self {
19         Self {
20             rx,
21             active: true,
22             buffer: MessageBuffer::new(),
23             interval: interval(INTERVAL),
24         }
25     }
26 }
47 }
```

In the `step()` method within the `select!` macro, add another branch that uses our `interval` field of type `Interval`, calling its `tick()` method.

The method waits for the interval to elapse. The countdown does not restart even if we did not wait for completion and called `tick()` again. This allows us to handle log events and react to the interval expiring simultaneously.

src/summarizer.rs

RUST

```
15 impl Summarizer {
33     async fn step(&mut self) {
34         select! {
35             message = self.rx.recv() => {
36                 self.handle_message(message);
37             }
38             _ = self.interval.tick() => {
39             }
40         }
41     }
49 }
```

To handle interval events, add a `handle_tick()` method and call it from the newly created branch in the `select!` macro.

In this method, we plan to check for new messages and send them for processing.

```
src/summarizer.rs RUST
15 impl Summarizer {
33     async fn step(&mut self) {
34         select! {
35             message = self.rx.recv() => {
36                 self.handle_message(message);
37             }
38             _ = self.interval.tick() => {
39                 self.handle_tick();
40             }
41         }
42     }
50     fn handle_tick(&mut self) {
51     }
52 }
```

To check for new messages in `MessageBuffer`, add a `has_new()` method that returns `bool`, so it can be used in `if` blocks.

In its implementation, call the `is_empty()` method provided by the `Vec` type; in our case, this is the `new_messages` field that holds new messages.

To know whether the buffer of new messages is not empty, add the negation operator `!` to the value returned by `is_empty()`.

`src/buffer.rs`

RUST

```
4  impl MessageBuffer {
13     pub fn has_new(&self) -> bool {
14         !self.new_messages.is_empty()
15     }
16 }
```

Use the newly added `has_new()` method on `MessageBuffer` within `handle_tick()` to run code when new messages arrive.

Add an `if` block that wraps the summary-generation block when the buffer has new messages—that is, when `has_new()` returns `true`.

```
src/summarizer.rs RUST
15 impl Summarizer {
50     fn handle_tick(&mut self) {
51         if self.buffer.has_new() {
52             }
53     }
54 }
```


Extend Summarization

Buffer

To generate a short message that describes the changes and errors mentioned in new log messages, add the `generate_summary()` method.

src/summarizer.rs

RUST

```
15 impl Summarizer {  
54     fn generate_summary(&mut self) {  
55     }  
56 }
```

In the `handle_tick()` method, call `generate_summary()` if there are new messages in the buffer.

It's also helpful to indicate in the logs that you are starting to generate the summary, so add an appropriate message using the `info!` macro from the `tracing` crate.

```
src/summarizer.rs RUST
15 impl Summarizer {
50     fn handle_tick(&mut self) {
51         if self.buffer.has_new() {
52             tracing::info!("Sending logs for summarization...");
53             self.generate_summary();
54         }
55     }
58 }
```

Now we need to add the `build_messages()` method, which will be responsible for constructing messages for the LLM.

In this case, these are different messages (not just from logs), and they are needed to prepare the context and obtain a response from the AI model.

You can also call the added `build_messages()` method directly from our `generate_summary()` method, thereby preparing the message for sending.

A separate method is needed here because LLM messages can have different, sometimes verbose types, and this is a fairly heavy structure that is convenient to isolate in its own method.

src/summarizer.rs

RUST

```
15 impl Summarizer {  
56     fn generate_summary(&mut self) {  
57         self.build_messages();  
58     }  
59     fn build_messages(&mut self) {  
60     }  
61 }
```

To build context, it helps to send not only new messages but also include a previously processed message in the context, making it easier for the model to understand what has been happening in the logs overall.

For such a buffer, we need to know its size—how many messages we want to store. So add a `DEPTH` constant of type `usize`.

The `usize` type is used to represent sizes in the system. Assign it a value; for example, start with `1024`. This is a fairly large window of log messages, but since these are prior logs, it will not make the summary large.

```
src/buffer.rs
```

```
RUST
```

```
1 const DEPTH: usize = 1024;
```

Now let's add another collection to our `MessageBuffer` for processed messages.

A regular vector won't do; we should use `VecDeque` to efficiently implement a sliding window. We don't want to store all messages, only a limited number of them.

Import this collection from the standard library's `collections` module and add a `processed_messages` field of this type, containing strings, to the `MessageBuffer` struct.

`VecDeque` itself is not size-limited, but we can enforce this later by adding an explicit check that drops extra messages on insertion.

Also, in the `new()` function, include the `processed_messages` field when constructing the instance.

For its value, create the collection with `with_capacity()`, passing our constant as the parameter. This will preallocate space for that number of elements.

src/buffer.rs

RUST

```
1 use std::collections::VecDeque;
3 pub struct MessageBuffer {
4     new_messages: Vec<String>,
5     processed_messages: VecDeque<String>,
6 }
7 impl MessageBuffer {
8     pub fn new() -> Self {
9         Self {
10             new_messages: Vec::new(),
11             processed_messages: VecDeque::with_capacity(DEPTH),
12         }
13     }
20 }
```

Add the `get_messages()` method, which we will use to retrieve slices of processed and new messages.

src/buffer.rs

RUST

```
7  impl MessageBuffer {  
20      pub fn get_messages(&mut self) {  
21      }  
22 }
```

Since `VecDeque` with processed messages (the `processed_messages` field) allows adding elements at both ends, to obtain a slice from such a collection you must call the `make_contiguous()` method, which turns the collection into a contiguous sequence.

This lets you get a slice, which is exactly what is returned as the result. Call this method and store the resulting slice in the `processed` variable.

src/buffer.rs

RUST

```
7  impl MessageBuffer {  
20      pub fn get_messages(&mut self) {  
21          let processed =  
          self.processed_messages.make_contiguous();  
22      }  
23  }
```

For new messages, since this is a regular vector, you can simply take a slice with the `&` operator and assign it to the `new` variable.

src/buffer.rs

RUST

```
7  impl MessageBuffer {
20      pub fn get_messages(&mut self) {
21          let processed =
self.processed_messages.make_contiguous();
22          let new = &self.new_messages;
23      }
24 }
```


We now have both slices for the processed and new messages, and we can return them as a pair by setting the method's return value accordingly.

That is, a tuple containing slices of strings—slices of log file entries (slice type `&[String]`).

```
src/buffer.rs RUST
7  impl MessageBuffer {
20  pub fn get_messages(&mut self) -> (&[String], &[String]) {
21      let processed =
22          self.processed_messages.make_contiguous();
23      let new = &self.new_messages;
24      (processed, new)
25  }
```

Construct Summary Messages

We will continue implementing the `build_messages()` method at the point where we called `get_messages()` to obtain two slices: decompose the resulting tuple into two variables, `processed` and `new`.

These contain, respectively, previously processed messages for which we already obtained a summary, and new messages whose processing should be prioritized by the model.

src/summarizer.rs

RUST

```
15 impl Summarizer {  
59     fn build_messages(&mut self) {  
60         let (processed, new) = self.buffer.get_messages();  
61     }  
62 }
```

To gauge how much information is being sent to generate the summary, you can add a debug message with the number of new and processed messages using the `debug!` macro from the `tracing` crate.

```
src/summarizer.rs RUST
15 impl Summarizer {
59     fn build_messages(&mut self) {
61         tracing::debug!(
62             "Generating summary for {} new messages (context: {}
processed)",
63             new.len(),
64             processed.len()
65         );
66     }
67 }
```

Build part of the context from the processed messages. Add an `if` block for this purpose in the `build_messages()` method.

As the condition of this block, check whether there are any messages in the `processed` slice using the `is_empty()` method.

Assign the result of the entire `if` block to the `processed_context` variable, and then we will construct the corresponding parts of the context in each branch of the conditional.

src/summarizer.rs

RUST

```
15 impl Summarizer {  
59     fn build_messages(&mut self) {  
66         let processed_context = if processed.is_empty() {  
67             } else {  
68             };  
69     }  
70 }
```

If the queue of previous messages is empty, we can simply create the string “No previous errors” by calling `to_string()`, which converts a string literal into a `String`, since that is exactly what the `format!` macro used in the other branch returns.

If we already have processed messages, we can use the slice method `join()` to combine the entries into text joined by the delimiter passed as the method’s argument. In our case, this will be `\n`.

And with the `format!` macro, we will create a complete description by adding the prefix “*Previous errors (for context only)*” to indicate that this part consists of messages that have already been processed.

Use the `{}` placeholder to append the tail of text with the processed messages, separated by the delimiter.

Now the `processed_context` variable holds the message context, whose value depends on whether we have processed messages or not.

src/summarizer.rs

RUST

```
15 impl Summarizer {
59     fn build_messages(&mut self) {
66         let processed_context = if processed.is_empty() {
67             "No previous errors.".to_string()
68         } else {
69             format!(
70                 "Previous errors (for context only):\n{}",
71                 processed.join("\n")
72             )
73         };
74     }
75 }
```

New log entries will always be present, since we already verified that this part of the buffer is not empty.

So also use the `format!` macro to create the second part of the context by adding text that indicates the intent “*New errors (focus on this)*”.

Also use the slice’s `join()` method, passing a newline separator to gather all new messages into a single block of text.

Use the `{}` placeholder to insert the collected messages at the end of the formatted string.

And store the entire result in the `new_errors` variable.

src/summarizer.rs

RUST

```
15 impl Summarizer {  
59     fn build_messages(&mut self) {  
74         let new_errors = format!("New errors (focus on these):  
    \n{}", new.join("\n"));  
75     }  
76 }
```

Now it's time to construct well-formed messages to send to the language model.

Import the `async-openai` crate by adding it to the `dependencies` section of the `Cargo.toml` configuration file.

It provides an OpenAI API client along with the necessary types.

Since functionality in different parts is hidden behind feature gates, you can simply enable them all with the `full` feature.

Cargo.toml

TOML

```
5 [dependencies]
6 anyhow = "1.0"
7 async-openai = { version = "0.31.1", features = ["full"] }
8 clap = { version = "4.5", features = ["derive", "env"] }
9 futures = "0.3.31"
10 regex = "1.12"
11 tokio = { version = "1.48", features = ["full"] }
12 tracing = "0.1"
13 tracing-subscriber = { version = "0.3", features = ["env-
  filter"] }
```

We need to import many different types from this crate, and it's convenient to do so by importing everything from the `types` module using `*`.

We will need the `chat` namespace now and `audio` later, but we won't have to import the latter.

```
src/summarizer.rs
```

```
RUST
```

```
3 use async_openai::types::*;
```


The first type we need from this crate is the system request arguments with the long name `ChatCompletionRequestSystemMessageArgs` from the `chat` module.

The type implements the `Default` trait, so call the `default` method to get an empty list of arguments.

This arguments type effectively acts as a builder for constructing a system message for the AI model. It therefore provides the `build()` method to construct the message.

Call it and save the result in the `system_message` variable.

src/summarizer.rs

RUST

```
16 impl Summarizer {  
60     fn build_messages(&mut self) {  
76         let system_message =  
           chat::ChatCompletionRequestSystemMessageArgs::default()  
77         .build();  
78     }  
79 }
```

When constructing messages, an error may occur, so we should account for it and return the `Result` from the `build_messages()` method.

Also use the `?` operator to handle the result returned by the `build()` method.

```
src/summarizer.rs RUST
16 impl Summarizer {
60     fn build_messages(&mut self) -> Result<()> {
76         let system_message =
            chat::ChatCompletionRequestSystemMessageArgs::default()
77             .build()?;
78         Ok(())
79     }
80 }
```

Define System Prompt

Now we need to add the main prompt that defines the agent's functionality.

To do this, create a `system.md` file in the `prompt` folder.

Include it in the code using the `include_str!` macro by specifying the path to the created file, and assign it to the `SYSTEM_PROMPT` constant.

```
src/summarizer.rs
```

```
RUST
```

```
10 const SYSTEM_PROMPT: &str = include_str!("../prompts/system.md");
```

Define the agent's role as a log analyzer, and specify its goal: to create a short 1–2 sentence summary.

Be sure to emphasize new messages.

`prompts/system.md`

MARKDOWN

```
1 You are a log analysis assistant. Generate a brief 1-2 sentence  
summary of the NEW errors only.
```

Although the summary is based on new log entries, you must note that older messages are also present and necessary for a full understanding of the situation. Indicate this.

`prompts/system.md`

MARKDOWN

- 1 You are a log analysis assistant. Generate a brief 1-2 sentence summary of the NEW errors only.
- 2 Use previous errors only for context to understand the situation.

Also note that the agent should describe the problem at a high level and avoid technical details, as we plan to convert this to audio and extracting technical details from an audio message will be difficult, and we only need the big picture of what is happening in the logs.

`prompts/system.md`

MARKDOWN

- 1 You are a log analysis assistant. Generate a brief 1-2 sentence summary of the NEW errors only.
- 2 Use previous errors only for context to understand the situation.
- 3 Focus on identifying the main problem without technical details.

It should also be noted that the text itself must be easy to understand, even for non-technical users. This ensures it is as concise and quick to grasp when listened to.

prompts/system.md

MARKDOWN

- 1 You are a log analysis assistant. Generate a brief 1-2 sentence summary of the NEW errors only.
- 2 Use previous errors only for context to understand the situation.
- 3 Focus on identifying the main problem without technical details.
- 4 Keep it simple and understandable for non-technical users.

Build Chat Messages

Specify the newly created prompt as the content for the system message using the `content()` method, passing the `SYSTEM_PROMPT` constant as the argument.

src/summarizer.rs

RUST

```
17 impl Summarizer {
61     fn build_messages(&mut self) -> Result<()> {
77         let system_message =
            chat::ChatCompletionRequestSystemMessageArgs::default()
78             .content(SYSTEM_PROMPT)
79             .build()?;
81     }
82 }
```


We have defined the agent's overall purpose; now let's add a user message—that is, the logs currently read by the application.

For this, we need the `ChatCompletionRequest` type from the `chat` module.

Use the `default()` method to create a builder for these arguments.

Then call the `build()` method to construct the message, handling any potential errors. Store the result in the `user_message` variable.

src/summarizer.rs

RUST

```
17 impl Summarizer {  
61     fn build_messages(&mut self) -> Result<()> {  
80         let user_message =  
            chat::ChatCompletionRequestUserMessageArgs::default()  
81         .build()?;  
83     }  
84 }
```

For the message itself, simply pass the processed logs stored in the `processed_context` variable.

And the new entries are already concatenated into a text representation and stored in the `new_errors` variable.

Combine all of this using the `format!` macro and pass it as an argument to the `content()` method, which we also used earlier to populate the system message.

src/summarizer.rs

RUST

```
17 impl Summarizer {  
61     fn build_messages(&mut self) -> Result<()> {  
80         let user_message =  
chat::ChatCompletionRequestUserMessageArgs::default()  
81         .content(format!("{}", processed_context,  
new_errors))  
82         .build()?;  
84     }  
85 }
```

Although we constructed different message types, all of them must be converted to the common `ChatCompletionRequestMessage` type.

And return the constructed set of messages as a `Vec` from the `build_messages()` function.

You can immediately use the `vec!` macro to create an empty message vector, assign it to the `messages` variable, and return it as the result.

```
src/summarizer.rs RUST
17 impl Summarizer {
61     fn build_messages(&mut self) ->
      Result<Vec<chat::ChatCompletionRequestMessage>> {
83         let messages = vec![
84             ];
85         Ok(messages)
86     }
87 }
```

`ChatCompletionRequestMessage` is an enum with different variants.

For a system message, use `System`, and for a user message, use the `User` variant. Use these variants to add the messages to the vector you created.

src/summarizer.rs

RUST

```
17 impl Summarizer {  
61     fn build_messages(&mut self) ->  
        Result<Vec<chat::ChatCompletionRequestMessage>> {  
83         let messages = vec![  
84             chat::ChatCompletionRequestMessage::System(system_message),  
85             chat::ChatCompletionRequestMessage::User(user_message),  
86         ];  
88     }  
89 }
```

Since `build_messages()` now returns a vector of messages wrapped in a `Result`, we also need to handle errors, so `generate_summary()` should return a `Result` type as well. Add it as the return type.

And return `Ok` as the successful result.

src/summarizer.rs

RUST

```
17 impl Summarizer {  
58     fn generate_summary(&mut self) -> Result<()> {  
59         self.build_messages()?;  
60         Ok(())  
61     }  
90 }
```

Move up the call stack, since `generate_summary()` now also returns a `Result`. On error, we must propagate it further from the `handle_tick` method. There, also declare the return type as `Result`, and return `Ok` for the success case.

You've seen a good example of how to handle errors in a convenient, simple, and elegant way. Just pass them upward as a `Result` value and try to avoid custom errors or any intermediate handling unless it's truly necessary.

It may seem that you might miss some important details, but even the generic `Error` from the `anyhow` crate is more than sufficient. If necessary, you can always add additional context to the returned result when it's an error.

src/summarizer.rs

RUST

```
17 impl Summarizer {
52     fn handle_tick(&mut self) -> Result<()> {
53         if self.buffer.has_new() {
54             tracing::info!("Sending logs for summarization...");
55             self.generate_summary()?;
56         }
57         Ok(())
58     }
91 }
```

Accordingly, `handle_tick()` can now also fail, so the `step()` method itself must return a `Result` as well. Add `Ok` at the end as the success value.

When calling your `handle_tick()`, use the question mark operator to handle the results returned by this method.

```
src/summarizer.rs RUST
17 impl Summarizer {
35     async fn step(&mut self) -> Result<()> {
36         select! {
37             message = self.rx.recv() => {
38                 self.handle_message(message);
39             }
40             _ = self.interval.tick() => {
41                 self.handle_tick()?;
42             }
43         }
44         Ok(())
45     }
92 }
```

In the `run()` method, we do not want to interrupt the agent if something goes wrong while constructing or sending the `summary`. Therefore, we will not propagate the error upward here—though we will technically keep that option—and will handle it instead.

To do this, add an `if let` construct and check whether the result is the `Err` variant. If it is, print the error to the console and let the loop continue without breaking.

`src/summarizer.rs`

RUST

```
17 impl Summarizer {
28     pub async fn run(mut self) -> Result<()> {
29         tracing::info!("Starting summarizer");
30         while self.active {
31             if let Err(err) = self.step().await {
32                 tracing::error!("Error during processing:
33 {err}");
34             }
35             Ok(())
36         }
94 }
```


Configure Completion Request

So far, we have only assembled a set of messages for the AI model. To turn this into a summary, we need to make a request to the model by passing these messages.

Add a new `request_completion()` method for this purpose, and have it accept as its argument a vector of messages of type `ChatCompletionRequestMessage`, exactly those we returned from the `build_messages()` method.

src/summarizer.rs

RUST

```
17 impl Summarizer {
94     fn request_completion(
95         &self,
96         messages: Vec<chat::ChatCompletionRequestMessage>,
97     ) {
98     }
99 }
```

Now we can pass messages from one method to another.

In the `generate_summary()` method, we called the `build_messages()` method. Store the return value in the `messages` variable, then call the `request_completion()` method, passing the message vector as an argument to this new method.

src/summarizer.rs

RUST

```
17 impl Summarizer {
62     fn generate_summary(&mut self) -> Result<()> {
63         let messages = self.build_messages()?;
64         self.request_completion(messages);
65         Ok(())
66     }
100 }
```

We will now use a builder to prepare a request for the LLM model. The `async-openai` crate provides the `CreateChatCompletionRequestArgs` type—this is a builder for creating a completion request to the model.

Call its `default()` method to create the builder, then call the `build()` method to construct the request.

src/summarizer.rs

RUST

```
17 impl Summarizer {
95     fn request_completion(
96         &self,
97         messages: Vec<chat::ChatCompletionRequestMessage>,
98     ) {
99         let request =
100             chat::CreateChatCompletionRequestArgs::default()
101                 .build();
102     }
```

As with the previous builders, the build can fail, so the `request_completion()` method must also return a `Result`.

The result returned by the `build()` method should be handled using the `?` operator, and on success we should return `Ok`.

src/summarizer.rs

RUST

```
17 impl Summarizer {
95     fn request_completion(
96         &self,
97         messages: Vec<chat::ChatCompletionRequestMessage>,
98     ) -> Result<()> {
99         let request =
100             chat::CreateChatCompletionRequestArgs::default()
101                 .build()?;
102         Ok(())
103     }
104 }
```

To send a completion request, you must specify the model to use. For this purpose, we will choose `gpt-4.1-mini`.

Save this value in the `MODEL` constant and use the `model()` method of the `ChatComp` builder, passing the value of `MODEL` as the parameter.

The `gpt-4.1-mini` model was chosen for a reason, as it supports fine-tuning: you can adjust parameters such as temperature and set limits.

src/summarizer.rs

RUST

```
11 const MODEL: &str = "gpt-4.1-mini";
18 impl Summarizer {
96     fn request_completion(
97         &self,
98         messages: Vec<chat::ChatCompletionRequestMessage>,
99     ) -> Result<()> {
100         let request =
chat::CreateChatCompletionRequestArgs::default()
101             .model(MODEL)
102             .build()?;
104     }
105 }
```

Add the message passed in the `messages` parameter to the request. To do this, use the `messages()` method of the same name and pass the received vector as the argument.

```
src/summarizer.rs RUST
18 impl Summarizer {
96     fn request_completion(
97         &self,
98         messages: Vec<chat::ChatCompletionRequestMessage>,
99     ) -> Result<()> {
100         let request =
101             chat::CreateChatCompletionRequestArgs::default()
102                 .model(MODEL)
103                 .messages(messages)
104                 .build()?;
105     }
106 }
```

You can also set the number of output tokens for this model using the `max_completion_tokens()` method.

As the method parameter, pass the number of tokens the model may generate; it will not exceed this during generation.

We will use the `TOKEN_LIMIT` constant as the value, declared separately with the `u32` type the method expects, and set to 150 tokens.

As you likely already know, a token is not the same as a word. It is the smallest unit of text processed by the model and is usually part of a word. It is impossible to know in advance how many words will result from a given number of tokens.

```
src/summarizer.rs RUST
12 const TOKENS_LIMIT: u32 = 150;
19 impl Summarizer {
97     fn request_completion(
98         &self,
99         messages: Vec<chat::ChatCompletionRequestMessage>,
100     ) -> Result<()> {
101         let request =
chat::CreateChatCompletionRequestArgs::default()
102             .model(MODEL)
103             .messages(messages)
104             .max_completion_tokens(TOKENS_LIMIT)
105             .build()?;
107     }
108 }
```

We can also use the `temperature()` method to set the temperature—the model’s level of creativity or strictness.

For our case, set a lower temperature; for example, a value of `0.3` makes the output less creative and more formal and strict.

Store this value in a `TEMPERATURE` constant of type `f32`, since the temperature is a floating-point number, and pass this constant as an argument to the `temperature()` method.

```
src/summarizer.rs RUST
13 const TEMPERATURE: f32 = 0.3;
20 impl Summarizer {
98     fn request_completion(
99         &self,
100         messages: Vec<chat::ChatCompletionRequestMessage>,
101     ) -> Result<()> {
102         let request =
chat::CreateChatCompletionRequestArgs::default()
103             .model(MODEL)
104             .messages(messages)
105             .max_completion_tokens(TOKENS_LIMIT)
106             .temperature(TEMPERATURE)
107             .build()?;
109     }
110 }
```


Configure OpenAI Client

We now have a constructed request for the AI model, and we can create and initialize an API client to send it.

First, we need the configuration provided by the `OpenAIConfig` struct from the `config` module of the `async-openai` crate.

Import it and create a default instance using the `new()` method, storing the result in the `config` variable.

```
src/summarizer.rs RUST
3 use async_openai::{config::OpenAIConfig, types::chat};
20 impl Summarizer {
21     pub fn new(
22         rx: mpsc::UnboundedReceiver<String>,
23     ) -> Self {
24         let config = OpenAIConfig::new();
25         Self {
26             rx,
27             active: true,
28             buffer: MessageBuffer::new(),
29             interval: interval(INTERVAL),
30         }
31     }
111 }
```

However, the API cannot be used without a key, so we need to extend the configuration with an API key for accessing the model, which can be set using the `with_api_key()` method by passing it as a parameter.

To make the `Summarizer` universal, add an `API_KEY` parameter to the `new()` function and use this parameter when calling the aforementioned method.

```
src/summarizer.rs RUST
20 impl Summarizer {
21     pub fn new(
22         api_key: String,
23         rx: mpsc::UnboundedReceiver<String>,
24     ) -> Self {
25         let config = OpenAIConfig::new().with_api_key(api_key);
26         Self {
27             rx,
28             active: true,
29             buffer: MessageBuffer::new(),
30             interval: interval(INTERVAL),
31         }
32     }
112 }
```

It's convenient to obtain the API key from the command-line arguments. To do this, add an `api_key` field of type `String` to the `Args` structure.

Have this argument support `short` (short form) and `long` (long form), specified in the `arg` attribute.

You can also allow this parameter to be populated from an environment variable by adding the `env` attribute with the variable name.

Specify `OPENAI_API_KEY` as the environment variable.

Now your program can accept the key for the OpenAI client; pass it from the parsed `args` structure as an argument to the `new()` function that creates the `Summarizer` structure.

`src/main.rs`

RUST

```
11 #[derive(Parser)]
12 struct Args {
13     #[arg(value_name = "LOG_FILE")]
14     log_file: PathBuf,
15     #[arg(short, long, env = "OPENAI_API_KEY")]
16     api_key: String,
17 }
18 #[tokio::main]
19 async fn main() -> Result<()> {
24     let mut tasks = Vec::new();
27     let summarizer = Summarizer::new(args.api_key, log_rx);
28     tasks.push(summarizer.run().boxed());
31     select_all(tasks).await.0?;
32     Ok(())
33 }
```

We now have the configuration we created in the `config` variable, suitable for creating an OpenAI API client.

Now import the `Client` type from the `async-openai` crate.

Then use the `with_config()` function to create a client instance and store it in the `client` variable.

Pass the configuration we created, stored in the `config` variable, as the argument to the `with_config()` function.

```
src/summarizer.rs RUST
3 use async_openai::{config::OpenAIConfig, types::*, Client};
20 impl Summarizer {
21     pub fn new(
22         api_key: String,
23         rx: mpsc::UnboundedReceiver<String>,
24     ) -> Self {
25         let config = OpenAIConfig::new().with_api_key(api_key);
26         let client = Client::with_config(config);
27         Self {
28             rx,
29             active: true,
30             buffer: MessageBuffer::new(),
31             interval: interval(INTERVAL),
32         }
33     }
113 }
```

Also add a `client` field of type `Client<OpenAIConfig>`, since the client's `new` requires you to specify the configuration as a type parameter.

Add this field to the `Summarizer` struct, and in the `new()` function initialize it with the client you created earlier and stored in the `client` variable.

```
src/summarizer.rs RUST
14 pub struct Summarizer {
15     rx: mpsc::UnboundedReceiver<String>,
16     active: bool,
17     buffer: MessageBuffer,
18     interval: Interval,
19     client: Client<OpenAIConfig>,
20 }
21 impl Summarizer {
22     pub fn new(
23         api_key: String,
24         rx: mpsc::UnboundedReceiver<String>,
25     ) -> Self {
26         let config = OpenAIConfig::new().with_api_key(api_key);
27         let client = Client::with_config(config);
28         Self {
29             rx,
30             active: true,
31             buffer: MessageBuffer::new(),
32             interval: interval(INTERVAL),
33             client,
34         }
35     }
115 }
```

Wire AI Summarization

Flow

Now we can use the client to interact with the AI model by accessing the `Client` field of the struct.

The `Client` type provides the `chat()` method, which returns an object for working with the chat completion API.

src/summarizer.rs

RUST

```
21  impl Summarizer {
103      fn request_completion(
104          &self,
105          messages: Vec<chat::ChatCompletionRequestMessage>,
106      ) -> Result<()> {
113          self.client.chat();
115      }
116 }
```

The Chat Completion API provides the `create()` method, which generates a model response from the request passed as an argument.

Pass the `request` we created as the parameter. Because the method is asynchronous, use `await` to wait for the result. However, the function must be marked `async` to use this operator.

Store the result in the `response` variable. Also remember to handle the returned value, since it is wrapped in a `Result`.

src/summarizer.rs

RUST

```
21 impl Summarizer {
103     async fn request_completion(
104         &self,
105         messages: Vec<chat::ChatCompletionRequestMessage>,
106     ) -> Result<> {
113         let response =
self.client.chat().create(request).await?;
115     }
116 }
```

We should now extract the summary from the response.

The model may return multiple answers, so the response includes a `choices` field, which is a vector.

Access this field and store the result in a new `summary` variable.

```
src/summarizer.rs RUST
21  impl Summarizer {
103      async fn request_completion(
104          &self,
105          messages: Vec<chat::ChatCompletionRequestMessage>,
106      ) -> Result<()> {
114          let summary = response
115              .choices;
117      }
118 }
```


Since `choices` is a `Vec` of `ChatChoice` elements, it may contain many items, but we only need the first one.

You can specify this number in the request, but since we didn't, we will have only a single option anyway.

To obtain the first element of the vector, we first turn it into an iterator with `into_iter()`, then extract the first item with `next()`. Because a vector does not guarantee that it contains at least one element, the result is returned wrapped in an `Option`.

`src/summarizer.rs`

RUST

```
21 impl Summarizer {
103     async fn request_completion(
104         &self,
105         messages: Vec<chat::ChatCompletionRequestMessage>,
106     ) -> Result<()> {
114         let summary = response
115             .choices
116             .into_iter()
117             .next();
119     }
120 }
```

Therefore, we use `and_then()` and pass a closure that, from this `ChatChoice`, accesses the `message` field, which is of type `ChatCompletionResponseMessage`, and returns the `content` field, which is also optional and thus wrapped in `Option`.

Thus, by combining `and_then()` with the optional content inside, we handle the nested `Option` in one go.

src/summarizer.rs

RUST

```
21 impl Summarizer {
103   async fn request_completion(
104       &self,
105       messages: Vec<chat::ChatCompletionRequestMessage>,
106   ) -> Result<()> {
114       let summary = response
115           .choices
116           .into_iter()
117           .next()
118           .and_then(|choice| choice.message.content);
120   }
121 }
```

We should always expect a value, so convert this `Option` into a `Result` using the `context()` method from the `Context` trait in the crate `anyhow`, and provide a message indicating that the request has no response, which will serve as the error message.

We can now propagate any error in the result using the `?` operator.

```
src/summarizer.rs RUST
2 use anyhow::{Context, Result};
21 impl Summarizer {
103     async fn request_completion(
104         &self,
105         messages: Vec<chat::ChatCompletionRequestMessage>,
106     ) -> Result<()> {
114         let summary = response
115             .choices
116             .into_iter()
117             .next()
118             .and_then(|choice| choice.message.content)
119             .context("No content in OpenAI response")?;
121     }
122 }
```

Now we can change the return type of the `request_completion()` method to `String` and, at the end of the method, replace the returned value with the `summary` variable.

```
src/summarizer.rs RUST
21  impl Summarizer {
103      async fn request_completion(
104          &self,
105          messages: Vec<chat::ChatCompletionRequestMessage>,
106      ) -> Result<String> {
120          Ok(summary)
121      }
122  }
```

Since the `request_completion()` method is asynchronous, you need to use the `await` operator to execute it. Therefore, make the `generate_summary()` method that calls it asynchronous as well using the `async` keyword.

```
src/summarizer.rs RUST
21 impl Summarizer {
70     async fn generate_summary(&mut self) -> Result<()> {
71         let messages = self.build_messages()?;
72         self.request_completion(messages).await?;
73         Ok(())
74     }
122 }
```

Let's propagate this asynchrony further up. Make the `handle_tick()` method asynchronous as well, and use the `await` operator at the point where it calls the `generate_summary()` method.

```
src/summarizer.rs RUST
21 impl Summarizer {
63     async fn handle_tick(&mut self) -> Result<()> {
64         if self.buffer.has_new() {
65             tracing::info!("Sending logs for summarization...");
66             self.generate_summary().await?;
67         }
68         Ok(())
69     }
122 }
```

Since the `step()` method was originally asynchronous, we only need to use the `await` operator when calling `handle_tick()`. This executes the entire chain of asynchronous methods.

```
src/summarizer.rs RUST
21 impl Summarizer {
45     async fn step(&mut self) -> Result<()> {
46         select! {
47             message = self.rx.recv() => {
48                 self.handle_message(message);
49             }
50             _ = self.interval.tick() => {
51                 self.handle_tick().await?;
52             }
53         }
54         Ok(())
55     }
122 }
```

Now we can store the returned value without calling the `request_completion()` method.

Add a `summary` variable, change the return type of `generate_summary()` to `String`, and use the created variable as the return value.

```
src/summarizer.rs RUST
21 impl Summarizer {
70     async fn generate_summary(&mut self) -> Result<String> {
71         let messages = self.build_messages()?;
72         let summary = self.request_completion(messages).await?;
73         Ok(summary)
74     }
122 }
```


Also store the `summary` returned by the `generate_summary()` method in a variable. You can log the summary using the `info` macro. This confirms that it was generated successfully.

```
src/summarizer.rs RUST
21 impl Summarizer {
63     async fn handle_tick(&mut self) -> Result<()> {
64         if self.buffer.has_new() {
65             tracing::info!("Sending logs for summarization...");
66             let summary = self.generate_summary().await?;
67             tracing::info!("Summary: {}", summary);
68         }
69         Ok(())
70     }
123 }
```

Manage Processed Messages

Now we need to improve the buffer. Add a `mark_as_processed()` method whose purpose is to move messages from the list of new ones to the list of processed ones.

src/buffer.rs

RUST

```
7  impl MessageBuffer {  
25      pub fn mark_as_processed(&mut self) {  
26          }  
27  }
```

With the `drain()` method we can obtain an iterator that removes elements. The method takes as its parameter a range specified by the `..` operator.

If you do not specify the start and end of the range, the entire array will be drained. Use this method by adding it to a `for` loop and assigning each element's value to the loop variable `message`.

src/buffer.rs

RUST

```
7  impl MessageBuffer {  
25      pub fn mark_as_processed(&mut self) {  
26          for message in self.new_messages.drain(..) {  
27              }  
28      }  
29  }
```

Messages retrieved in the loop from the new messages queue should be added to the processed list using the `push_back()` method, passing the extracted message as an argument.

src/buffer.rs

RUST

```
7  impl MessageBuffer {
25      pub fn mark_as_processed(&mut self) {
26          for message in self.new_messages.drain(..) {
27              self.processed_messages.push_back(message);
28          }
29      }
30 }
```

As you remember, we planned to store only a limited number of processed messages.

We already declared the `DEPTH` constant, which limits the number of messages, and we used it when initializing the vector size.

Now add an extra condition before adding a message to ensure that the queue length does not exceed this limit.

In other words, compare the value returned by the `len()` method with the limiting `DEPTH` constant.

If the limit is reached or exceeded, meaning our message queue is already full, use the `pop_front()` method to remove the oldest message.

This will discard the extra message.

src/buffer.rs

RUST

```
7 impl MessageBuffer {
25     pub fn mark_as_processed(&mut self) {
26         for message in self.new_messages.drain(..) {
27             if self.processed_messages.len() >= DEPTH {
28                 self.processed_messages.pop_front();
29             }
30             self.processed_messages.push_back(message);
31         }
32     }
33 }
```

We implemented the `mark_as_processed()` method, which moves new messages into the processed ones.

We will call this method after successfully generating a summary for the current set of processed and new log messages.

So call the `mark_as_processed()` method on the `buffer` field of the `Summarizer` struct in the `handle_tick()` method.

src/summarizer.rs

RUST

```
21 impl Summarizer {
63     async fn handle_tick(&mut self) -> Result<()> {
64         if self.buffer.has_new() {
65             tracing::info!("Sending logs for summarization...");
66             let summary = self.generate_summary().await?;
67             tracing::info!("Summary: {}", summary);
68             self.buffer.mark_as_processed();
69         }
70         Ok(())
71     }
124 }
```

Our minimal implementation is ready; it reads a file and generates a summary.

Let's test it. To do this, we need to set an OpenAI API key, which you can obtain on the OpenAI service management page.

After you set the environment variable, run the `cargo run` command, passing the name of the sample log file we created earlier (the `sample.log` file) as a parameter.

To pass parameters not to the `run` command but to our application itself, you need to use a double dash `--` as a separator between the command's parameters and our program's parameters.

```
. SH
$ export OPENAI_API_KEY="<your-key>"

. SH
$ cargo run -- sample.log
```

Ready for More?

You’ve already built a log companion that watches your files, pulls out important errors and warnings, and turns them into short, clear summaries. It keeps a rolling memory of past problems, focuses on the newest issues, and uses an AI prompt designed for simple, non-technical explanations that are easy to understand as audio.

In the full version of the book, you’ll hook those summaries up to text-to-speech, turning each update into MP3 audio and logging how much data you generate. You’ll add a dedicated audio worker that receives these audio chunks over channels, plays them through the system output using **Rodio**, and wire everything into your existing async task system so text and sound stay in sync.

By the end, you’ll have a fully automated incident “radio” that listens to your logs and speaks out what’s going wrong in real time.

Subscribe to Premium

- ✓ Full source code of any step
- ✓ Interactive web book
- ✓ Multiple languages
- ✓ Extra chapters
- ✓ Vibe-coding prompts!

visit knowledge.dev to learn more!