

# 1 On Permission Fatigue

2 I use [Kiro](#) at work and [Claude Code](#) on weekends. I also occasionally tinker with Cursor, Codex, Pi, and  
 3 others. The interfaces differ but these tools are all wrappers around models that are getting better at  
 4 an astounding speed. It is clear that we are not going back to writing code the way we did in the recent  
 5 past. It is also clear that babysitting these tools with Do you want <Tool> to do <X>? is already getting  
 6 in the way. As Claude Code's current documentation puts it:

7 *By default, Claude Code requests permission for actions that might modify your system: file writes, Bash*  
 8 *commands, MCP tools, etc. This is safe but tedious. After the tenth approval you're not really reviewing anymore,*  
 9 *you're just clicking through.*

10 – <https://code.claude.com/docs/en/best-practices#configure-permissions>

11 We cannot address *safe but tedious* with `--dangerously-skip-permissions`. Agents are tenacious and  
 12 even mild prodding can push them to seek unintended goals in [creative ways](#). If they have access to  
 13 sensitive data, can run arbitrary code, and can access remote endpoints, we have a problem. This  
 14 page shares how I use Apple Containers and Claude Code sandbox to contain this [lethal trifecta](#) in my  
 15 weekend projects.


## 16 My Claude Sandbox

17 I built [claude-sandbox](#) to automate this setup. It requires macOS Tahoe on Apple Silicon with the  
 18 [container](#) CLI installed. You can install it with Cargo:

```
19 cargo install --git https://github.com/aldrin/claude-sandbox.git
```

20 The [README](#) has the details, but the workflow is straightforward.


21 I go into a project I need AI assistance with and run `claude-sandbox init` to scaffold a `.claude-sandbox/`  
 22 folder with the container configuration. Then I run `claude-sandbox build` to build the container image  
 23 and `claude-sandbox run` to start a Claude Code session and work as usual. The build step installs the  
 24 latest Claude Code into the image and configures its sandbox settings. The only use for Claude on the  
 25 host-side is running `claude auth login` to populate the OAuth token. At runtime, `claude-sandbox run`  
 26 reads it from the host keychain and passes it to the Claude Code instance inside the container.



27 *Figure 1: The claude-sandbox workflow.*

28 To limit the data the agent can access, `claude-sandbox run` launches Claude Code inside an [Apple Con-](#)  
 29 [tainer](#). Each container runs in its own lightweight VM using `Virtualization.framework` on Apple Silicon,  
 30 so isolation happens at the VM boundary rather than through kernel namespaces. The container

- 31 mounts only the project directory I started it in. The agent has no path to any other file on my machine.  
 32 I pick what to share by picking where to run the tool.



33 *Figure 2: Isolation layers at runtime.*

- 34 To limit the impact of arbitrary code, `claude-sandbox` configures Claude Code's [Bash sandbox](#) inside  
 35 the container. With `autoAllowBashIfSandboxed` on, Bash commands run without approval prompts. The  
 36 `acceptEdits` permission mode lets the agent write files freely within the mounted directory. Under  
 37 the hood, [bubblewrap](#) restricts Bash commands to the mounted project directory and blocks network  
 38 access to domains not on the allowlist. Every Bash command and its child processes inherit these  
 39 restrictions, so if the agent goes off track, the damage stays contained.

- 40 To limit connections to arbitrary endpoints, all network traffic routes through local proxies. Anthropic's  
 41 [engineering post](#) describes the design: HTTP, HTTPS, and SOCKS traffic go through proxy servers on  
 42 localhost. DNS resolution and direct TCP connections to external hosts are blocked. When the agent  
 43 tries to reach a domain not on the allowlist, the request goes to Claude Code's permission system and  
 44 I see a prompt.

- 45 While the agent does its thing in the sandbox, the project directory is a bind mount, so changes appear  
 46 on my host filesystem in real time. I keep my own editor and git workflow. When the session ends, the  
 47 container goes away, but the work remains in my git workspace to review, refine, and publish. This page  
 48 was written in one such session.