



margo serve

Purpose

`serve` builds, watches, and serves a site from memory with live reload. It accepts zero or one directory or config path; the current directory is the default. A directory containing `site.yaml` uses that config. Any other directory uses default linked-site behavior.

Input and output

When `--port` is omitted, Margo selects an available port. `--open` asks the operating system to open the reported URL. Startup and successful build events go to stdout. Failed rebuilds go to stderr and keep the last successful snapshot. A later successful build replaces it and triggers browser reload.

Configured output is ignored by the watcher and is never written to disk by `serve`. For configured sites, conventional top-level sibling directories used for build outputs, reports, logs, and temporary artifacts—including `build/`—are ignored as well. The configured source tree and local assets declared by `site.yaml` remain watchable. The command always emits text diagnostics; it has no `--diagnostics` flag.

Examples

```
margo serve ./docs --host 127.0.0.1 --port 8080
```

Expected startup output follows this shape:

```
Margo development server (not for production)
Serving http://127.0.0.1:8080/
```

Failures and diagnostics

`serve.port_invalid` reports an explicit port outside `1` through `65535`.

`serve.input_unreadable` reports a missing input. `serve.input_invalid` reports a non-directory input that is not a YAML config. Watcher and listener failures use `serve.watch_*` or `serve.listen_*` codes. Fatal command failures exit `1`.

Limitations and care

This server has no TLS, authentication, authorization, or rate limiting. It is for development only. Loopback is the default. Binding beyond loopback prints a warning to stderr and exposes the preview to the reachable network.