LABS UTILITY

Intelligrit Labs | 15 commits | 1 primary builder

Turn text-heavy content into structured audio people can actually use.

page2pod takes web pages and local HTML, extracts chapters, caches audio generation by section, and outputs practical podcast files instead of one disposable text-to-speech dump.

15	1	H2	MIT
COMMITS	PRIMARY BUILDER	CHAPTER HEURISTIC	LICENSE

The Short Version

page2pod is a small tool with a strong Intelligrit fingerprint: start from real operator pain, keep the inputs raw, structure the output, and avoid regenerating work you already did. Instead of treating text-to-audio as a one-shot batch job, it treats a document as a durable artifact with chapters, hashes, cache state, and reusable output.

> *The interesting part is not that page2pod calls TTS. The interesting part is that it*

> *turns a page into an inspectable, repeatable media pipeline.*

What The Code And History Show

The commit history is short but clear. The project moved from a direct HTML-to-audio utility toward a more useful system: AI chapter extraction became the default, code block handling gained explicit modes, markdown input support was added, and caching remained central so only changed chapters need regeneration.

- Structured chapter extraction instead of monolithic output.
- Per-chapter caching for low-cost updates.
- JSON chapter index output for downstream players.
- Support for both URLs and local files.
- A clean command-line surface for automation.

Why It Matters

page2pod reflects several Intelligrit values directly: text-first workflows, accessible output, practical AI use, and durable intermediate files instead of opaque magic. The same design instinct matters in document-heavy internal systems and retrieval workflows.

Where It Fits

page2pod belongs in Labs as a good example of how Intelligrit builds bounded AI utilities that convert raw content into more usable forms without creating unnecessary platform drag.