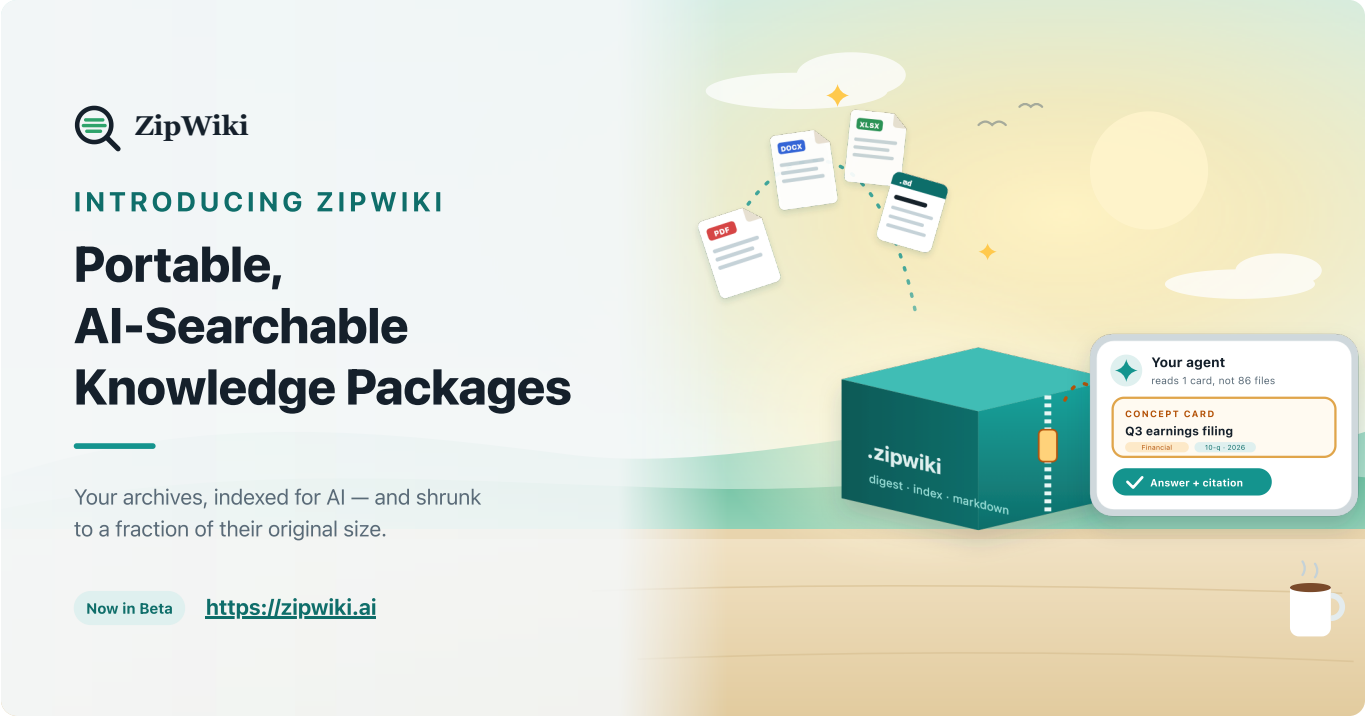



ZipWiki by NeoWare

## Why long-term archives need an AI upgrade

To an AI agent, a ZIP file is a blind box. It can see the filenames—and almost nothing else. To answer a question, the archive has to be unpacked and the agent has to chew through massive, heavy files from scratch every single time.

Every team accumulates folders of PDFs, Word documents, spreadsheets, and slide decks—session laws, quarterly earnings filings, clinical records, contracts, or reference books—and we traditionally pack them into ZIP files to save space. That works for humans. It fails for AI.

ZipWiki solves this problem. It is portable document memory with an embedded AI index—not another cloud drive, and not a complex database you have to maintain. When you pack a folder into a .zipwiki file, ZipWiki reads your documents once, pulls out the clean text, and writes a smart index right inside the ZIP container.

## What lives inside a ZipWiki package?

A .zipwiki is still a standard ZIP file that opens with normal unzipping tools on any computer, but organized specifically for AI:

- **An AI Digest (The Summary):** A quick overview of what is inside the archive, so an AI agent knows what is in the box before doing a deep dive.
- **The Smart Index:** Lightweight “concept cards” for each document. This acts like the index at the back of a textbook, helping an AI instantly locate key facts without reading every page.
- **Clean Text (Extracted Markdown):** AI models struggle with complex PDFs and binary files. ZipWiki converts documents into clean, structured text that AI can read rapidly.

- **The Originals (Optional):** Bulky original PDFs can stay bundled right next to the text, or they can be omitted to save space while keeping a verified web link to the original source.

## How AI reads a ZipWiki (Without getting overwhelmed)

Asking an AI a question about a massive archive shouldn't require loading the entire collection into memory at once.

ZipWiki is built for selective reading. When an AI interacts with a ZipWiki package, it follows a simple path:

- It reads the overall summary.
- It searches the smart index to identify which documents hold the answer.
- It streams only the specific text needed to answer your question.
- If it absolutely requires the original file, it knows exactly where to retrieve it.

## Shrinking knowledge to 1/100th the size

Because ZipWiki captures the text, tables, and headings an AI actually uses, bulky original files don't always need to travel with the archive.

Linking out to original files (such as state law portals or SEC databases) instead of packing them directly inside routinely shrinks an archive's size by upwards of 99%. That translates to roughly 1/100th of the original file size, while each answer still cites the concept card, the section it used, and the source link. When that original is fetched from the source link, ZipWiki can verify it against the recorded CRC-32. When full offline access is required, the original binaries can simply remain bundled inside.

## Trust the archive, not a hosted database

ZipWiki puts data back in your hands. There is no hosted index or vector store needed. You can share, sync, email, and back up .zipwiki archives like any standard ZIP file, confident that your knowledge remains instantly accessible to both humans and AI models.

Built by NeoWare, the team behind NeoZip. ZipWiki is now available in Beta. To explore the format, download the tools, and start packaging your archives, visit **zipwiki.ai**.

© NeoWare, Inc. · Introducing ZipWiki