

Mappit — Product & Technology Whitepaper

Version 1.9.0 · July 2026

The world's zeitgeist on a map. Every place has a story.

About this document. This is a plain-language explanation of what Mappit is, how it works under the hood, and how it was built — written for readers researching the product behind the **\$MAPPIT** token. It covers the **product and technology only**. It is **not** investment advice, not an offer or solicitation to buy or sell anything, and it makes **no claim about any token's price, value, or future performance**. Nothing here should be read as a promise. See the full disclaimer at the end.

1. Executive summary

Mappit turns the world's news into a living, interactive map. An always-on artificial-intelligence engine continuously scans what is happening in hundreds of cities — using **six independent, grounded sources** in parallel — and pins each story to the place it is about, on an interactive, time-aware map of the planet. Zoom into any city and you see what's happening there right now — grouped by topic, colored by mood, corroborated across multiple independent AI systems, and with every source link checked so the stories you read actually lead somewhere real. Scrub the calendar backwards and the map replays how the world looked on any past day, because the underlying record is permanent and never overwritten.

On top of that news layer, Mappit adds **live views of the physical**

planet — earthquakes, wildfires, air quality, and weather — and a **market layer** that plots stock movements at companies' real headquarters. A built-in AI guide named **Marco** answers questions about any place and cites its sources. Anyone can browse for free; a subscription unlocks the ability to create and share your own places, collections, and data layers.

The core idea is simple and durable: **place + time is the most natural way to organize what is happening in the world**, and the value compounds every day as the record grows. The map is the interface; the permanent, corroborated, geolocated corpus of world events beneath it is the substance.

2. The problem

The modern news environment is a firehose. It is:

- **Overwhelming and undifferentiated** — an endless scroll with no sense of *where* something is happening or how it connects to a place you care about.
- **Disconnected from time** — today's feed buries yesterday's, and there is no easy way to see what a place looked like on a specific day.
- **Hard to trust** — a single source can be wrong, biased, or fabricated, and readers rarely see whether a claim is independently corroborated.
- **Siloed from context** — the news about a city, its markets, its weather, its history, and its community all live in different apps.

Maps solve the "where." Archives solve the "when." Cross-checking solves the "is this real." Mappit is built to do all three at once, for the whole planet, continuously.

3. What Mappit is

Mappit is a web and mobile product centered on a single interactive map of the world. Its building blocks:

- **Locations** — cities and places, each a marker whose badge shows how many fresh stories are live there. Open one and you get a full card: stories grouped by category, community discussion, an AI-written overview of the place, and more.
- **Stories** — individual news items discovered by the AI engine, each geolocated to a place, dated, categorized, and tagged with how many independent sources corroborated it.
- **A time machine** — a date control that redraws the entire map as it looked on any chosen day or range. Nothing is ever deleted, so history is always replayable.
- **Overlays** — toggleable layers on top of the news: a mood heatmap, live planetary feeds (earthquakes, wildfires, air quality, weather), a stock-market layer, and a "this day in history" layer.
- **Marco** — an AI assistant that has "read" everything on the map and answers questions with clickable citations — including "what's trending today?", answered from the same daily brief subscribers receive by email.
- **Navigation** — street-address search with instant previews, reverse lookup ("what's here?") on any map point, and turn-by-turn driving directions between any two points, pins, or typed addresses.
- **Creation tools** — for subscribers: drop your own pins (places, events, street addresses), group them into shareable collections that compute into **drivable multi-stop routes**, import spreadsheets as map layers, and draw and measure regions.
- **Presence** — every member gets a personal **user pin**: put yourself on the map with one tap (Find Me), keep it current automatically as you move (Follow Me), and optionally breadcrumb your route into a shareable trail. Private by default; being findable is a choice.

Free to browse; a subscription unlocks creating and sharing. A **Daily Digest** email — an AI-written morning brief of the day's biggest stories plus each reader's subscribed places — goes to opted-in members.

4. How it works under the hood

This section explains the architecture at a conceptual level — enough to understand *why the product is hard to replicate* — without exposing the proprietary recipes, prompts, or operational details that make it run.

4.1 The AI news engine

The heart of Mappit is an autonomous engine that continuously sweeps a large roster of cities and asks: *what is happening here, and what happened in the wider world?* It does this using **six independent, state-of-the-art AI and data sources working in parallel** — leading models from Anthropic, Google, xAI, OpenAI, and Moonshot, each with live web access, plus signals from local online communities. All six now contribute to every sweep, so each story can be checked against the widest possible set of independent viewpoints. Each source runs across three rhythms: a sweep of individual cities, an hourly scan of the biggest stories on the whole planet, and a daily "this day in history" pass.

Four design choices make the output trustworthy and durable:

1. **Cross-model corroboration.** When two or more independent AI systems surface the same event, Mappit records that agreement as a **corroboration score**. A story confirmed by several independent models carries more weight than one seen by a single source — a built-in, automatic fact-cross-checking layer that no single-model system has, and one that gets sharper the more distinct providers agree.
2. **Deduplication.** The engine recognizes when different sources are describing the same event and merges them rather than double-counting, using both exact and fuzzy matching so near-duplicate headlines collapse into one corroborated story.
3. **Verified source links.** Grounded AI models have a well-known failure mode: they confidently cite article links that don't actually exist. Mappit measured this directly and found that, for the noisiest provider, **nearly half of its article links led nowhere** — dead pages or invented URLs. So Mappit now **checks every story's link before it reaches the map** and removes the fabricated ones,

with a recurring sweep that also cleans out links that have since gone dead. The result is a feed where a source link is a source link, not a guess — a genuine trust advantage over any system that simply prints whatever a model returns. The same principle now governs **story preview images**: rather than trusting a model to supply a picture, Mappit reads the *publisher's own* share image from the source article and confirms the picture actually loads before showing it — so a story's thumbnail is the real one the outlet chose, not an invented link. And where automated checks aren't enough, a curator can remove a bad story from the feed in two clicks, keeping the record clean by hand as well as by machine.

4. **An append-only, immutable corpus.** Stories are never overwritten or deleted. The record only grows. This is what makes the time machine possible — and it means the dataset beneath the product **compounds in value every single day**.

The engine runs on a disciplined schedule with automatic pacing, so busy places refresh often and quieter ones less so, and every run operates under a hard, configurable daily spending ceiling. That daily budget is now split into **two independent, separately-tracked pools** so that a hungry city crawl can never starve the essential world pulses that keep the map — and the Daily Digest — fresh. A **pulse pool** funds the hourly global scan and the daily "this day in history" pass, each with its own per-source budget; a **crawl pool** funds the city sweeps out of the remainder. The two pools have their own live meters and their own controls, so an operator can run a pulse or a crawl deliberately without one consuming the other's day. When the world's news demands it, the engine even **creates new places on the map** that didn't exist in its roster before — geolocating a dateline and adding the city automatically.

The engine is also relentlessly kept honest about the two failure modes that matter most here. The global pulse is now held to **strict, machine-readable output** — a source must return only stories with a real, live link, or nothing at all, never conversational prose — after models were observed drifting into apologetic explanations instead of results and returning empty sweeps; the long-standing rule against fabricated links is unchanged. And a source that stops working is pulled from the mix rather than paid for: when one model began failing every single call, it was simply disabled, so no budget is spent on sources that cannot

deliver.

Why six sources, and not just the cheapest one? Because cost and quality pull in different directions, and Mappit optimizes for the right thing: **cost per usable story**, not the raw price per story. A source can look cheap on paper yet be expensive in practice if many of its links are dead — you paid for stories you had to throw away. Mappit tracks the real, live cost of every source and blends them deliberately: the cheapest source widens coverage, a higher-integrity source anchors trust, and a newer low-cost source expands breadth further. This is a breadth-versus-quality-versus-cost balance that a single-provider product simply cannot strike.

4.2 The living map

Every story is automatically **categorized** against a consistent topic taxonomy (politics, business, culture, public safety, and more) and colored accordingly, so the map reads at a glance. Nearby places cluster together when you zoom out and separate as you zoom in, always keeping the busiest, most newsworthy places visible.

A **sentiment heatmap** overlay paints each place's mood — tense to upbeat — as a soft bloom of color sized by how much is happening there, so you can read the emotional temperature of an entire region in one look.

Two complementary views help you read the moment: a **trending** strip that ranks what matters most in a place (weighted by how many independent sources corroborate it), and a chronological "**Latest**" feed that surfaces headlines the instant they are discovered. Latest is ordered by *when Mappit actually found the story*, not by the publication dates the AI sources report — because those dates, like the links, can be unreliable — so "just in" really means just in.

For readers who want the whole planet at once rather than one place at a time, a top-level **Newsfeed** — reached straight from the main navigation, right under Home — is an infinite, strictly chronological river of story cards, newest first, running all the way down to the earliest stories Mappit ever ingested. Cards flow into responsive masonry

columns — one to four across depending on screen width — so the newest always sit top-left and older stories only extend the bottoms of the columns as you scroll, with no jarring reflow. Each card shows the publisher's own preview image when there is one and a button to fly straight to that story's place on the map. Every story card also now wears a small **provider-attribution mark** — the brand of the AI system, or the community, that surfaced it — so you can see at a glance where a story came from.

4.3 The time machine

Because the corpus is append-only, Mappit can reconstruct the map for any point in the past. Pick a date or a range and the stories, events, and activity from that window are what you see. This turns the product from a news feed into a **queryable historical record of the planet** — a fundamentally different and more valuable asset than a feed that forgets.

4.4 Marco — the AI guide

Marco is Mappit's built-in assistant. Ask it "what's happening in Lisbon?", "what happened here last month?", or "find community groups near Denver," and it answers from what is actually on the map — never from guesswork. Under the hood it uses **semantic search over the corpus** (a vector database that finds the most relevant stories, place overviews, and community content for your question) and then composes an answer that **cites its sources**; click a citation and the map flies to it. You can even pin a question to a **specific date or range**. If Marco doesn't have the answer in its corpus, it says so rather than inventing one — and it is designed to resist attempts to manipulate it through the content it reads.

4.5 Live planetary feeds and markets

Beyond the news, Mappit overlays the physical and economic world in real time:

- **Live planetary feeds** — recent **earthquakes**, active **wildfire**

hotspots, current **air quality**, and **weather**, sourced from authoritative public data providers and refreshed automatically. These are built on a reusable internal engine, so new live layers can be added quickly.

- **Markets** — an end-of-day stock layer that plots companies at their real headquarters and lets you watch them move green and red across any historical date, spanning multiple exchanges (including U.S. markets and the London Stock Exchange), each on its own trading calendar. A ticker tape rides the bottom of the map.
- **Media Streams** — live news broadcasts, radio and television, pinned to the city they broadcast from. A dedicated **Media overlay** (off by default) reveals them with their own glyphs — a satellite dish for TV, a radio tower for radio — and Audio/Video sub-filters. Click one and it plays in a **persistent player that floats in the corner of the screen and keeps playing as you browse the entire site**, so a station stays with you while you explore the map. Radio plays through a real audio element with a live frequency visualizer — bars that climb green to yellow to red — reading the actual signal where the station allows it and falling back to plain playback with a decorative meter otherwise, so the audio never cuts out; television plays as a live video embed. The starter set is seeded from an open, public radio directory (geolocated, news-tagged) plus a curated roster of round-the-clock news video channels, and subscribers can add their own by pasting a stream link and choosing audio or video.

Because these draw on public and low-cost data sources rather than the AI engine, they add breadth to the product without materially adding to its running cost.

4.6 Search that understands place and time

Search accepts topics, keywords, and place names — and also **dates**. Type a specific day, a range, or a whole month right in the query and Mappit narrows to stories from that window before matching your other terms. The same date-awareness works when you ask Marco a question.

4.7 Openness and discoverability

Mappit publishes machine-readable indexes of its public content so that search engines and AI assistants can discover and cite it. The corpus is built to be found and referenced — a growing, sourced, geolocated record that becomes part of how the wider web understands what is happening where.

5. Why it's defensible

Products are copied; **compounding datasets are not**. Mappit's durability comes from a handful of reinforcing advantages:

- **A permanent, growing corpus.** Every day of operation adds sourced, dated, geolocated, corroborated records that can never be regenerated retroactively by a competitor starting later. Time itself is the moat.
 - **Multi-model corroboration and verified links.** Trust is engineered in, not bolted on. Six independent sources cross-check each other, and every source link is validated before it reaches the map — two structural differentiators over any single-model approach that simply prints what a model returns.
 - **The data model, not the map.** The interactive map is the visible surface; the value is the structured record beneath it. That record powers search, the AI guide, the time machine, and future data products — and none of those depend on the specific map technology, which can evolve underneath without touching the data.
 - **A place-and-time interface people intuitively understand.** No training required: it's a map, and it's a clock.
-

6. The product experience

Anyone can browse for free — the map, the news, the history and live overlays, the market layer, place overviews, and a daily allowance of Marco questions. On a first visit the map centers on the user's location (or a sensible default) so it feels local immediately.

A subscription unlocks creation. Subscribers can drop their own **pins** (places, events, services, guides, meetups, street addresses, and live **media streams**), attach photos and media, group pins into shareable **collections** — and turn a collection into a **computed driving route** with total distance, time, and an optimized stop order. **Spreadsheet imports** become map layers; **draw and measure** covers regions. Paid tiers scale with how intensively someone tracks the world, and every tier is ad-free with a generous Marco allowance. The product is available on the **web and as a dedicated mobile experience** — not a shrunk desktop site, but a map-first interface designed for a phone, including full create-and-share on mobile.

Community and sharing run throughout: **shareable public profiles** (cover image, social links, your public map and posts), discussion on places and pins, following other contributors, a shared pin or profile acting as a session-long doorway to that creator's map, and rich link previews so any Mappit place, pin, collection, or profile unfurls with a branded card when shared in chat apps or social media. **Direct messages** now have first-class entry points on the desktop too — a mail icon in the header carrying an unread badge, and a left-navigation item. The **Daily Digest** email closes the loop each morning — reliably in the morning, now that both its once-a-day scheduling and its send-hour check read the same configured local time — with readership measured end to end.

7. How it was built

Mappit is notable not just for what it does but for **how efficiently it was created** — a signal of the engineering leverage behind the product.

It was built by a small operator using a **disciplined, AI-agent-coordinated development process**: a rigorous written architecture

specification first, then coordinated AI build agents executing it in waves with clear ownership, taking the system from an empty repository to a working, production-grade application in a matter of days rather than months. The result is a full modern stack — an interactive map front end, a robust server back end, a relational database with vector search for the AI features, and cloud media storage behind a global content network.

Crucially, the speed did not come at the expense of discipline. The system was engineered from day one with the safeguards you would expect from a much larger team:

- **Immutable data** — the corpus is append-only, so history is never corrupted.
- **Hard cost governance** — the AI engine operates under configurable spending ceilings and can be paused, cancelled, and resumed without losing work.
- **Full cost observability** — every AI operation is measured and attributed to its source in a live ledger, so spending is always visible and controllable, sources can be turned on or off individually, and pricing or model choices can be changed without redeploying. Mappit judges each source on cost per *usable* story — not just its headline price — so the provider mix stays honest.
- **Single-instance safety** — background engines coordinate so work is never duplicated, even across multiple servers.
- **Trust and safety** — corroboration scoring, content moderation tools, abuse rate-limiting, and guards against manipulation of the AI assistant.

This combination — extraordinary build velocity *plus* production discipline — is itself part of what makes the underlying operation efficient to run and extend.

8. Status and platform

Mappit is **live in production** and operating continuously. The news engine runs on an ongoing schedule, so the corpus grows every hour;

the map spans hundreds of cities with tens of thousands of corroborated stories and expands as the world's news demands. The product is available on the web and as a dedicated mobile experience, with a full subscription and billing system in place. Recent releases added the live planetary feeds and sentiment heatmap (1.2), the Global/Here/For You trending controls (1.5), shareable profiles, personal user pins, and the AI-written Daily Digest (1.6), the full navigation layer — address search, reverse lookup, directions, and computed collection routes (1.7), and the six-source feed engine with verified source links, cost-per-usable-story governance, and the chronological "Latest" view (1.8) — with 1.8.1 bringing all six sources fully online in every sweep and adding publisher-sourced, verified story preview images, and 1.9.0 adding live radio and television **Media Streams** on the map, a planet-wide chronological **Newsfeed**, and a **two-pool feed budget** that protects the essential world pulses from the city crawls.

9. Roadmap and vision

Mappit's direction follows from its thesis — *place and time, continuously, for everything*:

- **Deeper and broader coverage** — more cities, more languages, and more of the physical and economic world layered onto the map (additional markets, additional live feeds).
- **A next-generation map surface** — smoother, higher-fidelity rendering to support ever-richer layers and animations, evolving the visible surface without disturbing the underlying data.
- **Data products** — the compounding corpus (mood over time, attention by place and topic, market-versus-sentiment relationships) is a foundation for correlation-style datasets and discoveries that reach far beyond the map itself.
- **A citable, open record** — continuing to make the corpus discoverable so it becomes part of how the wider web and AI assistants understand what is happening where.

The through-line: the map is the beginning, not the product. The product is a permanent, growing, corroborated, geolocated record of the world —

and the many ways that record can be explored, queried, and built upon.

10. Disclaimer

This document is provided for **informational purposes only** to describe the Mappit product and its technology. It is **not** an offer to sell or a solicitation of an offer to buy any security, token, or other financial instrument, and it is **not** investment, financial, legal, or tax advice. It contains **no representation or promise regarding the price, value, utility, or future performance of any token, including \$MAPPIT**, and nothing herein should be relied upon in connection with any investment decision.

Statements about future plans and development are **forward-looking** and inherently uncertain; actual outcomes may differ materially, and Mappit undertakes no obligation to update them. Product features, availability, and status described here reflect a point in time and are subject to change. Any decision relating to \$MAPPIT or any other asset is made solely at the reader's own risk and should be based on independent research and professional advice.

Mappit — every place has a story.