

Mappit — Product & Technology Whitepaper

Version 2.9.6 "Frontier" · October 2026

Frontier geo-intelligence: 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. It covers the **product and technology only**. It is **not** investment, financial, or legal advice, and not an offer or solicitation to buy or sell anything. 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 news engine continuously reads what is happening in nearly two thousand places — **its own crawler**, which reads the world's publishers directly and supplies more than nine in ten new stories, working alongside **frontier AI from Google, OpenAI, Anthropic and xAI** in the jobs a crawler cannot do — 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, drawn only from publishers that clear a reliability bar and weighted toward the most trusted, marked with how many independent sources found it, 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 each day's record is kept rather than replaced.

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 — and since 2.8.0 that includes the lay of the land: the population, borders and statistics of any country, state or town of 500 people or more, answered from open reference data with the year and source of every figure. And the corpus is now also a **playing field**: three built-in games — a daily treasure hunt, a spot-the-fake quiz, and a multiplayer strategy game — are generated from the real stories of the day and the week, so playing them *is* reading the news. Anyone can browse for free; a subscription unlocks the ability to create and share your own places, collections, and data layers. And since 2.7.0 the map answers where people already ask: **Claude, ChatGPT and other AI assistants can read Mappit directly** through its MCP server — a briefing on any place, the storylines behind it, the businesses nearby — free and without an account, with every answer linking back to the publishers and to the map.

Mappit is also **frontier geo-intelligence**. Beside the news sit eight **Layer Apps**, a reference desk for the planet: 112 views that colour the world's countries and America's states, counties and congressional districts by the economy, population, energy, health, conflict and the 2026 US elections, from open data whose licences allow commercial reuse, and play them back through decades like a film. They draw the places that matter to those numbers, too: every nuclear reactor in the world and the daily status of America's, the data centers mapped in OpenStreetMap, and the frontier AI campuses as Epoch AI tracks them being built, each a click away from how far it stands from you. Every figure names its source. Readers ask the map in plain English and it sets itself up, and since 2.9.6 any view can carry an **AI analysis**: Claude Sonnet, thinking hard, reads the figures, how they moved, the trend and the news in the same place, and writes what it sees, readable by everyone.

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, sourced, 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 found by Mappit's crawler or its AI search, each geolocated to a place, dated, categorized, carrying its publisher's reliability rating, and tagged with how many independent sources found it.
- **A time machine** — a date control that redraws the entire map as it looked on any chosen day or range. Stories are kept rather than replaced, 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, a "this day in history" layer, and national and state borders.

- **Layer Apps (map apps)** — eight sets of statistical views (World Economy, People, US Economy, Health, Environment, Conflict & Crisis, US Elections and News Atlas) that colour countries, US states, counties or congressional districts by a figure from open data, rank them, and play the years, months or weeks back like a film; draw sites such as every nuclear reactor and the world's data centers and AI campuses, with how far each is from you; and, for Plus supporters, write an AI analysis of any view that every reader can read. Every figure carries its source, period and licence.
- **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; since 2.8.0, questions about populations, borders and statistics for any place, answered from open reference data with the year and source of every figure; and, since 2.9.0, the map apps' figures and the 2026 US races, with a button that shows the answer on the map.
- **Mappit in your AI assistant** — one address, <https://mappit.ai/mcp>, lets Claude, ChatGPT and other assistants that speak the open Model Context Protocol read the map: briefings on any place, country or the whole world, the full story lists, storylines with Mappit's analysis, the Daily Digest and the Week in Review, businesses and landmarks, the live layers, the map apps' statistics and directions. Free, no account, read-only; every story links to its publisher and to its place on the map.
- **Navigation** — street-address search with instant previews, reverse lookup ("what's here?") on any map point, and turn-by-turn directions by car, on foot or by bike 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.
- **Games** — three news-native games generated daily and weekly from the live corpus: **Where's Marco** (a riddle trail across real

stories), **Bullshit** (spot the one doctored summary among four — three real, one with a single fact bent), and **Conquest** (a strategy game whose board is the week's news cities and whose economy is driven by how hot each storyline is, played against friends or against Marco, Mappit's AI general). Every clue and question links back to its source story, and every province to its city's storyline.

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 news engine

The heart of Mappit is an autonomous engine that continuously reads the world's newsrooms and asks: *what is happening in each place, and what is happening in the wider world?* Two routes feed it. **Marco the crawler**, built on what the first two months of AI queries taught the platform about who publishes the news where, reads those publishers directly at no search cost and supplies the overwhelming majority of new stories — since 2.6.0 including the places no big publisher covers, which it reads through their own local newsrooms. **Frontier AI from Google, OpenAI, Anthropic and xAI** does what a page's markup cannot. Models from Google and OpenAI — joined since 2.7.0 by Anthropic's Claude and xAI's Grok — run the scan of the biggest stories on the whole planet with live web search; Google's and OpenAI's models run the daily "this day in history" pass; and Google's Gemini also names each storyline and — once per place — finds the local newsrooms the crawler should read. A small OpenAI model files every crawled story under its topic, reads its mood and, where the crawler could only guess, says where the story happened; OpenAI embeddings power Marco's retrieval, the Threads clustering and the semantic corroboration pass. Anthropic's Claude powers Marco and writes the digests, the analyses

and the games. The crawler is the engine's volume and its reach; the AI systems are its judgment and its view of the whole world. A frontier model's find and a crawled article of the same event corroborate each other exactly as two independent sources should.

Five design choices make the output trustworthy and durable:

1. **Cross-source corroboration.** When independent sources — Mappit's crawler reading a publisher, and a frontier model's live search — surface the same event, Mappit records that agreement as a **corroboration score**, and the story strips weigh it heavily. It is a signal that more than one route found the story, not a check of every fact in it.
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. One wire is one story: when the same syndicated article appears on many publishers' sites, the first copy stands, and when copies merge, the best-rated publisher's link is the one that survives.
3. **Verified sources.** 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 **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 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 a real share image that a publisher of that story chose, and since 2.9.1 a frontier model's new story is kept only with a picture from the publisher's own page, the wires included: readers trust a story with its own picture more. 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. Further gates stand in front of the map: every publisher carries a **reliability rating** drawn from

Wikipedia's perennial-sources consensus, plus a preferred tier of the major US and UK wires, networks and papers of record — those lead the crawl and the story strips, and unreliable and deprecated publishers are never crawled; an operator-kept **publisher blacklist** refuses a listed domain at ingest and sweeps out anything already stored; and **Google Safe Browsing** — every article host is looked up in the same lists behind Google's transparency report before its stories are written, and a daily sweep re-checks every host on the map at least weekly, so a publisher flagged for malware or phishing is blacklisted automatically. Since 2.7.0 the reliability bar is a gate, not just a weight: a story from a publisher rated unreliable or deprecated, or one a curator has set aside, is refused at the door, and 5,470 such stories already stored were removed — quality over quantity, for readers and for the search engines that judge a site by its links. (A first pass also swept up newsrooms the crawler merely couldn't read, KING 5 among them; their stories were restored from the backup the next day.)

4. **The primary route to the news: Marco the crawler.** Asking frontier models to search was the expensive half of the engine — in early September, seventy percent of the feed's daily spend was search fees, not tokens. So as of 2.4.0 the corpus feeds itself: **Marco the crawler** walks the front pages of the publishers on its roster (every publisher in the archive plus those it has found on its own, tiered by how much each contributes and rated for reliability, so the biggest and most trusted are read every run, mid-sized ones every eight to sixteen hours, the long tail monthly), follows their links a few levels down, and reads each article's own markup — its share tags, structured data and dateline — with no AI search at all; the only per-story model call is a small one that files each story under its topic and reads its mood, and a story it cannot place with confidence is not kept. (The one exception is a single grounded AI question per place, asked once, to find the newsrooms of a place no news feed could name.) A page becomes a story only if it passes a quality gate at least as strict as a vendor's find: a real title and link, a recent publication date, a description, a share image that actually loads, a known language, and a place — read from the article's own geo tags or dateline, from a major city named in the headline, or from the publisher's home city when the article names none (never for a wire copy,

which must place itself). Since 2.7.0 those last two are checked: the same small model, asked independently where the story happened, moves it there — a Kyiv market strike that a mention of Moscow had filed under Moscow now sits in Kyiv — and a story that happened nowhere in particular, such as celebrity news or a recipe, is no longer pinned on the publisher's town. The stories placed that way before the check existed are being re-checked the same way. Those verdicts can also correct the homes themselves: a pass moves a New York paper that had been homed in a Westchester village to New York, and takes from a national outlet with no home town the one it was given. The first pass is the founder's, and every such change can be undone. And places that share a name are now told apart by the state or province a source names: Washington the state and Washington, D.C., had been trading stories — the White House and the Supreme Court filed in the state, and the state's stories, and other towns' named Washington, filed in the capital — and the misfiled stories are moved back once the 2.8.0 build starts. Crawled stories carry Marco's mark instead of a vendor logo, and they **corroborate** with the frontier models' finds in both directions. The crawler is polite (it honors every publisher's robots.txt rules and paces itself), breadth-first (every planned publisher has its front page read early, with deeper reads interleaved so stories flow from the first minute), and self-pruning (a publisher that yields nothing run after run is set aside, and one whose front page keeps refusing it is retried weekly rather than every run). Its first production run put 49 stories on the map in fifteen minutes with no model call at all; today the only per-story cost is the small model call that files each story under its topic — about a quarter of a dollar per thousand stories kept. As of 2.5.2 the crawler also finds its own publishers — a **scout** asks a free news feed who covers each place and puts those publishers on the roster with the place as their home, so a town of a thousand people gets its local paper read without a vendor search — and a **headline lane** beside the big run now laps the top eighty publishers every half hour, the most reliable first, so the big sites stay fresh while the long tail is walked. Since 2.6.0 a third, **location lane** reads, every hour, the places short of their day's crawled stories through their own newsrooms, the neediest first — which is what let the frontier models' scheduled place searches be switched off.

5. **A permanent, dated record.** A new day's stories are added to the record rather than replacing yesterday's, so any past day can be replayed and the dataset beneath the product **compounds in value every single day**. Stories leave the map only when they fail the engine's own checks — a dead or fabricated link, an unsafe, blacklisted or low-quality publisher, a duplicate of a story already there — and corrections (a better place, topic or picture) are made in place.

The engine runs on a disciplined schedule with automatic pacing, so busy places refresh often and quieter ones less so. The crawler is paced by a daily story budget spread across the day, and its headline and location lanes count against it. The frontier-model runs work under a hard, configurable daily dollar ceiling, split into **two independent, separately-tracked pools** so that one kind of work can never starve the other: a **pulse pool** funds the global scan and the daily "this day in history" pass that keep the map — and the Daily Digest — fresh, and a **crawl pool** funds the place searches an operator starts by hand (the scheduled ones were handed to the crawler in 2.6.0). The two pools have their own live meters and their own controls, and the month's AI spend, frontier models and crawler together, is tracked against an operating band. 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 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 failed every call it was given, it was simply disabled, so no budget is spent on sources that cannot deliver.

Why several 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 keeps only what earns its place: its

own crawler supplies the volume for about a quarter of a dollar per thousand stories kept, frontier models with live search cover the world pulse and the history pass, and the other vendors stay wired for on-demand checks. When the frontier models' place searches turned out to be buying stories for places the crawler already covered — 86% of their spend in the week measured — they were switched off.

4.2 Where the news comes from — and what Mappit deliberately doesn't use

Mappit's news comes from **publishers' own pages**, read directly by its crawler, and from **frontier models' live web search** on the global and history pulses. Three other well-known sources of news data have been tried or considered and are deliberately *not* used as story sources.

RSS — a directory, not a source. Mappit reads RSS in exactly one place. The crawler's scout asks Bing News, through its public RSS search feed, which publishers are reporting on a place, and keeps only the answer's *publishers*: it never takes the feed's links, headlines or text, and the crawler then reads each publisher's own site. The scout checks the feed's robots.txt first, as the crawler checks a publisher's. Google News closes its feeds to automated readers there, and Google's terms forbid automated access against those rules, so since September 2026 the scout no longer asks Google; most new places are now staffed by the single grounded AI question described above. Everywhere else, RSS is set aside on purpose. The crawler reads front pages and article markup, not feeds, and its link filter skips feed, RSS, Atom and sitemap links, because following them turns a front-page crawl into an archive walk. The reasons are practical:

- **The page is needed anyway.** A story passes Mappit's gate only on what its article page declares — a publication date, a share image that loads, a dateline or geo tags, a language — so a feed item is only a pointer to the page the crawler already reads.
- **The feeds that would matter most are gone or closed.** Reuters ended its public feeds in 2020, and its robots.txt prohibits automated collection; AP's robots.txt bars its feed paths; CNN's feeds stopped updating; USA Today publishes none; and a few publishers, The New York Times among them, offer feeds while refusing automated reads of the articles they point to.
- **Where a feed is open, the front page usually is too**, and it

carries the same stories with the markup the gate needs.

- **Feed terms** commonly permit a headline and a link, not the reuse of a publisher's text.

The one feed format on Mappit's roadmap is the **news sitemap**: several major publishers that close their pages to crawlers publish Google News sitemaps that their robots.txt allows, and reading those is the sanctioned route to them. Spoofing a browser past a bot wall is not.

GDELT — built, measured, and left off. The GDELT Project's DOC 2.0 API is free, needs no key, and returns real article links from a global news firehose, so Mappit built a provider for it in July 2026 as a possible free substitute for AI place searches. It never went into rotation, for four reasons:

- **Throttling.** The API allows roughly one request every five seconds per address, and a live test in September drew rate-limit refusals even at eight to sixteen seconds apart. At that pace one pass over nearly two thousand places takes four to eight hours, and does not reliably finish.
- **Precision.** GDELT's geographic endpoint has been retired, so the only way to ask it about a place is a full-text search for the place's name. That returns articles that *mention* a city rather than articles *datelined* there — the opposite of what a map needs.
- **Thin records.** Each result is a headline, a link, a domain and a time, with no summary, so every item would still need its article page read before it could pass Mappit's gate. That is exactly what the crawler already does, from the publisher itself.
- **Redundancy.** Since 2.6.0 the crawler's location lane delivers what GDELT was meant to deliver — free, per-place news with real links — from the places' own newsrooms.

The provider remains in the code, switched off; it never put a story on the map.

Reddit — where the name came from. Mappit began as the idea of a *geolocated Reddit*: the communities Reddit organizes around cities and regions, put on a map — which is where the name comes from. The first design paired AI search with those local communities, and the first build seeded four hundred place-based subreddits beside five hundred world cities. Reddit never became a source. In Mappit's first week Reddit's bot wall blocked both its development and production servers, and API

access turned out to require developer registration; the launch corpus was built without it, and in the meantime the product found its real subject — sourced, dated reporting from publishers, placed where it happened. Community posts are not that kind of article, so Reddit is not used and not planned. The provider remains in the code, switched off, and it never published a story.

4.3 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. A companion **activity** mode answers a different question — not how the world *feels* but where it is *loud* — weighting the same bloom by sheer volume of stories and discussion.

The map itself is a GPU-rendered vector canvas with **six ways to look at the world**: a warm daylight street map, a luminous night mode, satellite imagery, a hybrid of imagery with readable labels floated over it (since 2.9.1 the view a first visit opens on), a terrain view, and a grayscale "data ink" minimal mode that gets out of the way of the data. Relief is always present — real elevation shading — and can be lifted into **true 3D terrain**, with buildings rising at street zoom, atmospheric sky at a tilted camera, and optional topographic contour lines. The planet can be viewed as a **globe** rather than a flat sheet. Selecting a place arrives cinematically, with a tilted, slowly drifting camera rather than a jump cut. Every one of these is a preference, not a requirement: readers who ask their device for reduced motion get a still, instant map, and the whole visual layer is operator-controlled — any mode can be switched off centrally without shipping new software.

Since 2.8.0 a **Borders** layer draws national and state or provincial lines over any of the six views and on the globe — beneath the place names, in colours chosen for each view, with contested borders such as Kashmir and Crimea dashed — and each level can be switched on or off alone; since 2.9.0 both are on by default. It restyles the boundaries

already in Mappit's self-hosted map, so it added nothing to serve. Zoomed far out, state lines appear for some countries first — the US and Canada from the start, the French, Italian and Spanish regions a little closer in — and for every country at a regional zoom.

Because the map is a single canvas rather than a page of images, any view can be **shared as a picture** — the current camera, layers and all, composited with attribution and handed straight to a phone's share sheet.

And the camera itself has learned to move with meaning. A computed route — a delivery round, a road trip, a walking tour — can be **flown**: the camera rides the route low and tilted while the line draws itself just ahead, at the viewer's choice of speed. Live location tracking gets the same treatment: following your own movement now reads like a **chase shot**, the camera leaning into your direction of travel while your trail glows behind you, brightest where you just stepped. And the archive can be **replayed**: on a computer, press Rewind and the map plays itself back a day at a time — stories, pins, mood and markets stepping together through any window of the record — turning the time machine from a dial you set into a film you watch. Every one of these honors reduced-motion preferences and can be switched off centrally.

Two complementary views help you read the moment: a **trending** strip that ranks what matters most — weighted by how many independent sources corroborate a story, by what the strip is for, and by the publisher's reliability — 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 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 wears a

small **provider-attribution mark** — the logo of the AI vendor whose search surfaced it, or Marco's compass for a story the crawler read — so you can see at a glance where a story came from.

A place's own panel, and the river of its stories behind it, reads in **your language and the place's**: Barcelona in Spanish, Catalan and English, Montreal in English and French, Washington in English. Stories in other languages wait one click away, named by language, so a story in a far-off language filed under a place no longer greets a reader who came for its local news.

4.4 The time machine

Because every story is kept with its date rather than replaced, 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.

4.5 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, with citations. 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**. When the map has nothing on a question, Marco falls back on a stronger model's general knowledge — saying so when the question is about Mappit or something on the map — and those answers carry no citations. It is designed to resist attempts to manipulate it through the content it reads. Marco also knows the games are Mappit's own: it answers rules questions from the guide, and from a Conquest match page it holds that table's live brief — seats, treasures, who holds what, the latest dispatch — and coaches from a stated method: restate the position, name the biggest threat and the best opportunity, propose a concrete turn of orders with reasons, and never guess hidden orders or predict dice. And Marco knows the way to Mappit from other assistants: ask it "how do I add Mappit to

Claude?" and it gives the address, the steps for your assistant and a few questions to try. And since 2.8.0 Marco knows geography (§4.11): ask "how many people live in Nampa, Idaho?" or "which states border Texas?" and it answers from open reference data, with the year and source of every figure. Since 2.9.0 it also reads the map apps (§4.12): ask "what's Ohio's unemployment rate?", "which counties in Texas have the highest poverty?" or "what do the polls say in the Georgia Senate race?" and it answers with the figure, its period and its source, and adds a **Show on map** button that opens the view it read.

4.6 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 visualizer never silences a station; 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 — local newsrooms in US and

Canadian cities and national channels worldwide — that Mappit checks every half hour, playing whatever each channel is actually broadcasting and taking a channel off the map while it is off the air. 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.

4.7 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. Suggestions start at two letters, and a one- or two-letter term matches a whole word — "AI" finds stories about AI, not every headline with "said" in it — and every place in the results carries a **Newsfeed** button that opens all of that place's stories directly. Since 2.9.1 a long query, such as a whole Where's Marco riddle pasted in, runs on its eight strongest words, shown as chips you can remove or add back, and story searches answer in a fraction of the time they took.

4.8 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. Since 2.7.0 assistants don't have to wait for an index: they can ask Mappit directly (§4.10). Since 2.9.0 the page a search engine reads for a place or a storyline carries Mappit's own analysis of it, the index written for AI assistants lists the most recently active storylines that have one, and every map app view has its own link, each app has its own share preview and crawlable page, and each app is listed in the sitemap. Since 2.9.1 a shared map app view previews as a small map of itself, with its legend, its top regions and its source, and a shared Threads page as its top four storylines. Since 2.9.3 search engines also get a page for every map app view, carrying its latest figures and source; for every 2026 Senate, governor and House race; for every state's figure

for every statistic; and for every country's conflict map, each titled the way people search ("Ohio governor race 2026", "Inflation in California", "Ukraine conflict map"), and a place's page is titled as its local news ("Chicago, IL news today"). The search engines that take instant notices (Bing among them) are told the moment a storyline gets a new analysis or a view's figures change.

4.9 Threads — the news, grouped into storylines

Individual stories are the atoms; **Threads** are the molecules. A Thread is a living *event* — "*War in Iran*", "*Heat Waves*", "*France and Spain Wildfires*" — assembled automatically by grouping related stories into one running storyline you can follow as it develops. Threads are the natural counterpart to the map (which answers *where*) and the trending strip (which answers *what's hot right now*): they answer *what's the ongoing story, and how did it get here*.

The mechanism is what makes it cheap and durable. Every story is already turned into a mathematical "fingerprint" (an embedding) when it's ingested — the same fingerprints that let Marco search the corpus. Threads reuse those fingerprints to measure which stories are *about the same thing*, so the grouping costs no extra AI work; the new spend is a small request to name each cluster with a human-readable label and one-line summary (and, for the ten biggest storylines, a written analysis, refreshed every four hours while new stories arrive). Each Thread remembers its own fingerprint (the average of its members'), so as fresh stories arrive they join the nearest existing storyline, genuinely new events start new Threads, and near-duplicates merge — meaning a story that runs for a week stays **one** Thread that grows, rather than fragmenting into a new one every hour.

Because Threads are first-class, they thread through the rest of the product: you can **search** a storyline by name ("War in Iran" takes you straight to it), **Marco** reads each Thread's rollup and answers "what's the latest on X", a shared Thread link previews on social media as a **rich montage** of its member photos, and the feed can be **filtered by date**. It turns a firehose of individual headlines into a legible, followable set of stories — at almost no marginal cost, because it's built on infrastructure the platform already runs. Since 2.9.1 Threads list the storylines a reader can actually read: a thread appears when enough of its stories are in the reader's language, the rest are a toggle away, and a

storyline's page shows its stories in the reader's language first.

4.10 Mappit in your AI assistant

More and more, people ask an AI assistant before they open a website. Since **2.7.0 "Ask Mappit"**, Mappit answers there too. It runs an **MCP server** — MCP, the Model Context Protocol, is the open standard AI assistants use to reach outside tools — at one address, **<https://mappit.ai/mcp>**. Add that address to **Claude, ChatGPT, Claude Code, Cursor or VS Code**, and the assistant can read Mappit's map for you:

- *"What's happening in Seattle?"* — a briefing: the top stories as Mappit ranks them, the storylines running through the city, and live conditions.
- *"What's the latest on the war in Iran, and what does Mappit's analysis say?"* — a storyline, its stories and Mappit's own analysis.
- *"Who covers local news in Boise?"* — a place's profile, including the newsrooms that reported on it in the last month.
- *"Find ramen near Pike Place Market", "Any earthquakes near Tokyo this week?", "What happened on this day in London?"* — businesses and landmarks, and the live layers.
- *"Rank US states by overdose deaths", "What map apps does Mappit have?"* — since 2.9.0, the map apps' statistics: any figure for a region, or regions ranked by any indicator, each with its period, source and licence and a link that opens the view on the map.
- *"Walking directions from Pike Place Market to the Space Needle"* — since 2.9.2, directions by car, on foot or by bike between towns, businesses and landmarks, street addresses or coordinates, with stops on the way, and a link that opens the route on Mappit's map.

It is **free, needs no account, and only reads** — it cannot change anything on Mappit. Every story comes back with two links, the publisher's own article and the story on Mappit's map, and every answer is dated, so the assistant can say how fresh the news is and credit the newsrooms that reported it. The answers are the map's own: the same stories, ranked the same way, with Mappit's analyses labeled as Mappit's. Places are understood the way people say them ("Portland, ME", "Washington State", "UK"), and when several places share a name, the assistant is told which one answered and which others exist.

The server is built to be fair and observable: each question costs a few units from a generous daily allowance, and every call is logged — the tool, the place asked about, how long it took, whether it worked — so Mappit's operators can see which questions it answers well and **which places people ask about that it doesn't cover yet**. Mappit sees only what an assistant sends its tools, such as the place asked about — never the conversation, the user's other prompts or their files.

Why it matters: the record Mappit builds becomes a source the assistants cite, with links back to Mappit — distribution through the tools people already use, rather than a site they have to remember to visit. The first session with Claude showed a second benefit: an assistant reading the whole map at speed is a tireless reviewer, and its first findings — a story filed in the wrong city, history events pinned on the wrong town, a confusing field — were fixed within hours. Since 2.9.3 an agent can find the server before it connects, through a server card, links in the home page's headers and an API catalog, and the server was listed in its first directories; an assistant that sends a malformed request is told what to correct. Next come more directory listings; sign-in, so an assistant can save pins, follow places and ask Marco for a member; and the map itself drawn inside the conversation.

4.11 The lay of the land

Ask about a place and you often want more than its news: how many people live there, what it borders, how its economy compares. Mappit follows the news in nearly 2,000 places, and people ask about all the others too. Since **2.8.0 "Lay of the Land"**, Marco answers those questions from a reference library of **open geography data** that sits beside the news:

- **Every country and territory (252), every state, province and region (3,865), every county and district (47,643) and every town of 500 people or more** — 235,917 places, most of which have no news on Mappit yet.
- **Their names in about 30 languages** — some 945,000 of them — so a place is recognised the way people write it.
- **1.56 million dated statistics**: populations through the years, and 30 World Bank indicators for every country from 1990 to 2025, GDP per person among them.
- **Land borders with their lengths, and the languages each**

country speaks.

Every figure carries its **source, and its year** wherever the source gives one, because the sources disagree. For American places the US Census's 2025 estimates come first; elsewhere, the most recent dated figure in Wikidata; a figure with no date is marked as undated. A city's own population and its wider urban area's are kept apart, and Marco says which one it is giving. The sources are open, and the user guide credits them:

- GeoNames — CC BY 4.0
- The World Bank's World Development Indicators — CC BY 4.0
- The US Census Bureau — public domain
- Wikidata — CC0
- Natural Earth — public domain
- Unicode CLDR — the Unicode License
- The CIA World Factbook, final 2026 edition — public domain

When a question names a real place, Marco is handed a short, sourced fact card for it before answering, and for rankings, borders and distances it can look things up itself: *"Which country in Africa has the highest GDP per capita?"*, *"Which states border Texas?"*, *"How far is Lyon from Geneva?"* Places that share a name are settled the way a person would settle them: "Paris, Texas" is the Texas city, "Tbilisi, Georgia" the country and "Atlanta, Georgia" the state, and "TN" is Tennessee. Game words no longer get in the way, so "what's the capital of Moldova?" is a geography question, not a Conquest one. Questions about the news, or about Mappit itself, are answered as before.

Until now, Marco handed the questions the map could not answer — 38% of all its questions — to a larger model's general knowledge, and a population came back without a year or a source. In testing with real questions it now answers from the data — Nampa, Idaho: 120,384 people (2025, US Census); Texas's eight neighbours, with the length of each border; Lyon to Geneva, 112 km in a straight line — at about half a cent a question. The library is looked up directly rather than searched by AI, so it adds almost nothing to Mappit's running cost.

The same release draws national and state **borders** on the map (§4.3). Together, the library and the borders are the groundwork for the **map apps** (§4.12).

4.12 Map apps

The geography library answers a question about one place. Since **2.9.0 "Almanac"**, the map can show the answer for every place at once. A **map app** colours each country, US state, US county or congressional district by a figure — 4,406 regions in all, with the congressional districts drawn as they were in 2024 and as they are for 2026 — and plays the years, months or weeks back like a film. Each app is a set of **views** over one subject. Open one from the Layers menu or from a link, and a panel beside the map (a bottom sheet on a phone) switches views and levels, plays the periods, ranks the regions highest first, names the biggest movers, and opens any region's card with its line over time. There are eight apps, with 112 views between them:

- **World Economy** — GDP per person, growth, inflation, unemployment, inequality, poverty, government debt, currencies, interest rates and stock markets, country by country, most of them since 1990; and since 2.9.4 how much fertilizer each country's farms use.
- **People** — population, density, growth, births per woman, life expectancy, age, urban share and migration, with population and its growth for US states and counties too.
- **US Economy** — unemployment (states since 1976, counties since 1990), jobless claims, income, cost of living, home prices, poverty and household income; since 2.9.2 the household budget: inflation, gas and electricity prices, mortgage rates, new homes permitted, foreclosures, and personal bankruptcies county by county; and since 2.9.3 what households owe: debt per person, car loans, credit cards, student loans and mortgages, the share of each that is seriously late, and medical debt in collections, by state and county; and since 2.9.4 diesel prices, what farms spend on fertilizer per acre, and job growth, by state every month and by county every quarter.
- **Health** — COVID, flu and RSV hospital admissions, overdose deaths, measles and the uninsured in the US; life expectancy, child mortality, doctors, vaccination and health spending worldwide; and the World Health Organization's outbreak reports as points on the map.
- **Environment** (since 2.9.4) — temperature against normal and average temperature, country by country every month since 1950;

carbon dioxide per person and fine-particle air pollution worldwide; since 2.9.6 how each country makes its electricity (renewable, low-carbon, solar, wind, nuclear, coal and gas shares since 1985) and how much carbon its grid gives off; every nuclear reactor in the world over nuclear's share, from its first year to its closing, and the daily power level of each US reactor; every data center mapped in OpenStreetMap; the frontier AI campuses as Epoch AI tracks them, growing quarter by quarter, alone or over power prices and grid carbon; US temperature and rainfall against normal by state since 1895 and by county since 2015; the days of unhealthy air in each county with a monitor since 1980; and drought.

- **Conflict & Crisis** — conflict events replayed as points sized by their deaths, conflict deaths by country since 1989, refugees and the internally displaced, and US disaster declarations and drought.
- **US Elections** — every 2026 Senate, governor and House race, with the candidates' fundraising and Mappit's own average of the polls (for the House since 2.9.2); presidential and Senate margins by state since 1976, presidential margins by county since 2000 and the 2024 House to replay; since 2.9.5 the early vote, every state's ballots so far against its 2022 turnout from the University of Florida's Election Lab (used with its permission) and North Carolina's by county, with the early voters' party registration; and results on election night.
- **News Atlas** — Mappit's own reading of the world: how many stories each region makes, their mood, and the share about conflict, the economy or crime.

The figures come from **48 providers** that load open data on their own schedules: the World Bank; the US Bureau of Labor Statistics, Department of Labor, Bureau of Economic Analysis, Energy Information Administration, Census Bureau, Department of Agriculture, National Oceanic and Atmospheric Administration, Environmental Protection Agency, federal courts, Centers for Disease Control and Prevention, Federal Emergency Management Agency and Federal Election Commission; the Federal Housing Finance Agency; the Federal Reserve Bank of New York and the Urban Institute; the Bank for International Settlements and the OECD; exchange rates from the European Central Bank and other central banks; Our World in Data, for Copernicus's temperature record and the world's electricity mix; the World Health

Organization; the Uppsala Conflict Data Program; UNHCR; the US Drought Monitor; VoteHub's polls; the University of Florida's Election Lab and North Carolina's elections board; the MIT Election Data and Science Lab; Global Energy Monitor, the Nuclear Regulatory Commission, OpenStreetMap's contributors and Epoch AI, for reactors and data centers; and Mappit's own news, counted by region and day for the last 13 months. Only sources whose terms allow reuse by a commercial site are used, so some well-known ones — ACLED, the IMF and FRED among them — are not used directly (a few World Bank series, such as government debt, carry figures the IMF compiled, under the World Bank's open licence). Every figure carries its source, period and licence in the legend, the panel and Marco's answers; provisional and estimated values say so, a region without data is plain grey, and a view whose data has not arrived yet is not offered. As first loaded on 9 October 2026, the apps hold 3.85 million stored statistics for 73 of their 80 indicators (the three news indicators are counted from Mappit's own stories; the three 2026 results wait for election night, and one example source ships switched off), 431,627 conflict events since 1989, 2,766 outbreak reports, 5,278 US disaster declarations, and 506 races for 2026 — 35 for the Senate, including special elections in Ohio and Florida, 36 for governor and all 435 House seats — with 2,480 candidates.

Replay. Each period is coloured from its own figures alone, so its colours compare the regions within that period and the map looks the same however a reader reaches it: on load, by a link, by scrubbing, playing or rewinding. Since 2.9.3 a trend line over the timeline shows the change over time, each period's average across the regions with the current period marked, and a replay ends on the map it opened with.

Elections. Party colours — blue for Democrats, red for Republicans — appear only in the US Elections margin views; every other view, the news coverage of the elections included, uses neutral colour scales. A race card lists each party's likely nominee rather than every candidate who filed. On election night, 3 November, the results feeds of North Carolina, Colorado, Iowa and Georgia are read out of sight first, and their results are shown only once they have been checked.

Ask the map. A question typed into the panel — *"where is unemployment highest?", "show me poverty by county in Texas", "how did Pennsylvania vote in 2020?"* — switches the map to the view that

answers it, with a one-line answer worked out from the data. A small model, Claude Haiku 5.5, only chooses the view; Mappit checks every choice against the catalogue and computes the answer itself, so no number comes from the model. Asking runs under a budget of 50 cents a day, with a daily limit for each reader.

Sites. Since 2.9.6 some views draw places that open and close rather than events: a reactor stands from its first year to its last, an AI campus grows from a ring while it is built to a dot sized by the power it draws. Playing the view replays them by year, or by quarter for the young data center build-out, and each site's card says how far it is from the reader.

Analysis. Since 2.9.6 each app's panel has an **Analysis** tab. A Plus supporter presses Generate, and Claude Sonnet, thinking hard, reads the view as it stands at the period on screen (and for the region selected): the figures, the biggest movers, the trend, the region's other figures, the sites on the map and the news Mappit placed there, and writes what the map shows, what changed, why it might look that way, how the news bears on it and what to watch. Every number comes from the figures it was given. Everyone can read the analysis once written; it can be rewritten four hours later; each supporter writes five a day; and readers without Plus see the button and the way in.

Marco and the assistants. Marco looks figures up, ranks regions, follows a trend and describes a race, with a **Show on map** button under the answer (§4.5). Short briefs on the latest statistics, each race, each country's month of conflict and the apps themselves join the corpus it searches — 1,169 of them at the start, refreshed every six hours and re-indexed only when their figures change. Assistants connected through MCP (§4.10) get three more tools: the list of apps, any figure for a region, and regions ranked by any indicator.

AI-gathered figures. Some figures are published only as a table on an official web page. For those, Mappit fetches the page itself and a small model reads the table; every number must appear on the page exactly as written and pass range and change checks before it is stored, a new source runs out of sight for 14 days before its figures are shown, and the figures say "AI-gathered" wherever they appear. The spend is capped at \$10 a month; the one example source defined at launch was replaced in 2.9.2 by the same figures from the publisher's own data service, so no AI-gathered source runs today.

Apart from Ask the map, the analyses (a few cents each, capped at \$10 a day) and the AI-gathered figures, the apps run on free open data that is looked up directly rather than searched by AI, so they add little to Mappit's running cost. Anyone can browse them, with no account.

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 records — each marked with how many independent sources found it — that a competitor starting later could not easily reconstruct. Time itself is the moat.
- **Cross-source corroboration and verified links.** Trust is engineered in, not bolted on. Independent routes to the news — Mappit's own crawler reading thousands of publishers directly, and frontier-model live search — cross-check each other, every publisher carries a reliability rating, and every source link is validated before it reaches the map: structural differentiators over any 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.
- **Distribution through the assistants people already use.** Since 2.7.0 any AI assistant that speaks the Model Context Protocol can read Mappit's record directly, with links back to Mappit and to the publishers. The corpus becomes something assistants cite, not only a site people have to remember to visit — and every question they ask shows where coverage should grow next.
- **The news, joined to the lay of the land.** Since 2.8.0 Mappit holds both sides of an unusual join: a curated, placed, dated record of the news, and a sourced, dated reference of the world's

places — their populations, economies, borders and languages. Statistics sites tend to have the numbers without the news, and news products the news without the geography. Since 2.9.0 the two share one map: the map apps (\$4.12) paint official statistics and Mappit's own count of each region's news over the same regions, and Marco reads both. The features designed next (\$9) — region dossiers, a measure of which places the news overlooks — need both.

6. The product experience

The front door. Since 2.9.1 the welcome page presents Mappit as one world read three ways: visualize the news on the **map**, take a sip from the firehose in the **Newsfeed**, and see the storylines at a glance in **Threads**. It plays each, lets the newest stories and storylines arrive live as you read, replays fifty years of US unemployment state by state from the map apps, and opens straight into any of them. Since 2.9.4 a signed-out visitor is offered the free Daily Digest first — an email box in the hero, the world's news and their own city's every morning — with the map a quiet link beside it, because with a door into the map in every section the page had signed up 0.6% of the visitors from the ads, against 1.4% for the page before it.

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

Or don't open the site at all. Add <https://mappit.ai/mcp> to Claude, ChatGPT or a coding assistant and ask it what's happening anywhere: it answers from Mappit's map, credits and links the publishers, and needs no account. Marco explains how to connect each assistant.

Readers keep it running. Mappit carries no advertising, so nothing on it is tuned for outrage, and since 2.9.5 Mappit Plus is presented as what it is: a contribution, about a coffee a month, from the readers who keep the map alive, with its perks offered as thank-yous. Free members are asked once in each Daily Digest and in a short letter twice a month,

which Claude drafts and an administrator approves before it goes out. **A subscription also 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, each 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.

The games make the news playable. Three games run on the corpus itself, each an engine for reading the news without noticing:

- **Where's Marco** is a daily treasure hunt: each morning the platform selects a geographically spread trail of well-corroborated stories from the last day and writes five riddles, each pointing at a city through its news. A city that was an answer in the last two weeks sits out — at most one returns in a hunt, unless the day's news is too thin to fill the trail — so the trail keeps moving. Players deduce, fly the map there, and claim finds; each solved clue reveals the real story it was written from. The sixth and final find is Marco himself — the assistant's actual map pin moves to the day's finale. Scoring rewards completion and speed; hints (the clue's country, then a one-line description and a 30 km ring on the map) cost points; equal scores go to the faster run. A validator checks every generated riddle before publication — including a whole-word anti-spoiler check that the riddle never names its own answer.
- **Bullshit** is a weekly spot-the-fake in seven daily rounds: four story summaries — three real, one carrying exactly one mutated fact (an actor, a settled number, a place, or an outcome). All four summaries are rewritten in one uniform news-brief voice so the forgery cannot be detected by tone — only by knowing the news. One minute on the clock; the reveal names the exact mutation and links the real storyline. Sunday's round plays against history rather than the week's news. Generation deliberately refuses to mutate figures still rising (casualty counts, spreading totals) — the

summaries lag live events, and a "correct" answer must never mark a well-informed reader wrong.

- **Conquest** is a 2–5 player strategy game on a board that regenerates weekly: the week's most-covered news cities become provinces (their territories computed geometrically around them), and **the news is the weather** — a province whose storyline is trending pays up to double yields, and three daily event cards translate real stories into rule effects ("grid strain in a heat wave: energy yields –25% today"; on a day a real cyberattack makes the news, a cyber card lets a single hack break through). Players submit up to five orders a turn at first — move, attack, missile, influence, sabotage, cyber, assassinate, build, buy, politick — which lock on submission and resolve simultaneously the moment every player is in, as many turns a day as a table can play — and a one-week war never waits more than a day: a general who hasn't reported within 24 hours of a turn opening sits that turn out, so nobody is rolled over at midnight while they sleep, and an epic waits for everyone. Planes let an attack reach two borders out, over land or sea, and since 2.8.0, after the players' feedback, they cost more and need a scientist in the empire. A hacker's cyberattack works in two steps — the first probes a rival's network, the second breaches it and wipes one of their orders for the turn before anything else resolves — and only one scientist, hacker or politician counts in each province, so specialists spread across an empire rather than stack. A table starts when it fills or when its seated generals choose to (two or more), and starting positions are dealt fairly on the board's own borders — capitals at least three borders apart, each with a similar share of the board. Combat is dice-based with **provably fair rolls**: each turn's random seed is committed as a hash before orders and revealed after resolution, so any player can recompute every die. An AI war correspondent files a narrative dispatch each turn, naming whose province is whose, backed by a complete mechanical ledger that names each province's holder too, and a replay control walks the match back and forth, turn by turn. Weekly matches score on Friday; "blitz" tables play a full campaign in one sitting. Since 2.9.5, "Season Two", a new war's board wraps east to west, so no start is a corner, and each general's view opens with them in the middle; every seven turns is a week of the war, and each week

prices, recruits, orders and income grow, the cheap things fastest, so a hoard loses value; missiles strike an asset three borders away; satellites are bought for attack or for defense; assassins, the dearest unit, reach any province and need two attempts to kill a scientist, hacker or spy (one on a day an assassination makes the news); and generals are told, by a push or an email, when the table is waiting only on them. When a war ends, its dispatches open with the war in review, the whole war retold above a table of every general's points with the arithmetic shown. Since 2.9.6 a table can **reserve a seat for Marco**, Mappit's AI general: two seats with him make a solo war that starts at once, and at a bigger table he is a wildcard. He plays as a Cadet, a Captain or a General, or adaptively, starting from how a player's last war against him went and easing off when he leads, sharpening when he trails. At a table of three or more he can take a player's side: he never strikes them and can't win the war himself. He reads the same board every player sees, never their orders or the dice, his orders are checked against the game's rules before they lock, and when his model can't answer he plays a fixed playbook, so no war ever waits on him.

The games cross-pollinate — finding Marco's trail pays Conquest resources, calling the day's fake pays energy — and they share the platform's cost discipline: all generation runs under its own hard daily AI budget ceiling (a dollar a day across the three games), fails soft (a game that can't generate skips its day rather than degrading the app), and links every generated element to its source as a required field, not decoration. New generals get a guided **walkthrough** — the board, the turn and its thirteen-phase order of resolution, units, assets, dice worked through six examples, and a scripted four-turn practice match played by locking in prepared orders and reading the ledger that results.

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** 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:

- **Permanent, dated data** — each day's stories are added to the record rather than written over it, and stories are removed only when they fail the engine's own link, safety or duplicate checks.
- **Hard cost governance** — the frontier-model runs operate under configurable daily ceilings, the crawler under a daily story budget, and the month is tracked against an operating band; everything can be paused, cancelled, and resumed without losing work.
- **Full cost observability** — every AI operation is measured and attributed to its source in live ledgers, 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, publisher reliability

ratings, 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 crawler reads publishers around the clock, so the corpus grows every hour (since it was tuned on 5 October, the crawler alone adds 14,000 to 20,000 stories a day): as of 14 October 2026 the map holds more than 2,000 places and nearly 485,000 stories, grouped into more than 3,000 storylines, and it expands as the world's news demands. Beside them, the map apps hold 3.85 million statistics from open data. The product is available on the web, as a dedicated mobile experience and, since 2.7.0, inside AI assistants, with a full subscription and billing system in place.

The foundation (1.x, July 2026). 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), with every story's link checked on the way in from 1.7.1; the multi-vendor feed engine with cost-per-usable-story governance, dead links deleted rather than hidden, and the chronological "Latest" view (1.8), with 1.8.1 bringing OpenAI into every sweep and adding publisher-sourced, verified story preview images; 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 sweeps (1.9.0); **restored cross-source corroboration** (now maintained automatically by a semantic pass), a **daily geography-less news brief** so Marco answers "today's headlines" without a place named, an expanded Media Streams roster (**NPR to 81 U.S. stations, BBC to 48**, plus Reuters, CNBC and Newsmax on video — all three since retired), **hardened Daily Digest delivery** so recipients are never silently dropped, a **public Digest archive**, and **shareable routes** (1.9.1); and **Threads**, a top-level

feed that groups the news into emergent, persistent trending storylines ("War in Iran", "Heat Waves") by clustering the story embeddings the platform already builds — searchable, readable by Marco, and shareable as rich composite cards — plus **date filters** on the Newsfeed and Threads and a **Follow Me / Find Me overhaul** (1.9.2).

The map rebuilt (2.0–2.2, August 2026). **Version 2.0.0 "Atlas"** moved the canvas to **MapLibre GL** — GPU-rendered vector tiles with the **entire planet self-hosted** on Cloudflare's edge (no third-party map keys, no rate limits, no per-view costs), two hand-tuned basemap themes (a warm daylight "Streets" and a luminous night-mode "Dusk"), buttery fractional zoom, and drawing/measuring tools rebuilt natively for the new engine — with every existing layer, panel, and API contract unchanged. **Version 2.0.1 "Places"** put a business directory under the map: the full OpenStreetMap storefront layer (~15 million points — restaurants, hotels, pharmacies, shops) imported once and searched by name, surfaced as a **Places** tab in global search and read at retrieval time by **Marco** so questions about a named place answer with real coordinates, address, hours, and website — a deliberate lookup rather than RAG, at zero embedding cost. **Version 2.1.0 "Analysis"** added a News & Analysis layer: a curator can generate an AI analysis of any story thread (what's happening, why it matters, the connections across it, what to watch), shown on a new **Analysis** tab and publicly shareable; and each week a **"Week in Review"** analyzes the biggest storylines and ties them together in a cross-cutting meta-analysis, published to a new top-level **Analysis** page and emailed as a **weekly digest** alongside the daily one. **Version 2.2.0 "Cinema"** completed the map rebuild Atlas began, in two acts. The glow-up: six ways to view the planet (streets, night, satellite, hybrid, terrain, and a minimal data-ink mode) behind a basemap picker, real elevation relief with an optional 3D terrain view, 3D buildings, an opt-in globe projection, cinematic camera movement when you select a place, live hazard layers that visibly breathe at a scale matched to their severity, story-arc lines drawn between the cities a storyline touches, an activity heatmap alongside the mood one, and one-tap sharing of any view as a picture. Then the showpieces: **route fly-throughs** (any computed route becomes a low, tilted camera ride with the line drawing itself ahead), a **chase-shot follow camera** for live location tracking with a glowing trail, and **Rewind** — the archive played back day by day as a film. All of it honors **reduced-motion** preferences, every element is switchable from the operator console without a software release — and

with the rebuild proven, the previous map engine was **removed entirely**, leaving a single, GPU-rendered canvas.

The games (2.3, August–September 2026). Version 2.3.0 "Arcade" (August 2026) turned the corpus into a playing field: the three games above — Where's Marco, Bullshit, and Conquest — shipped as a suite with shared leaderboards (daily, season, all-time), browsable past editions, an in-product rulebook, touch-first controls, and an operator console card governing generation and spend; the whole suite was designed, built, deployed, and founder-playtested inside a single week, with the launch-week lessons (anti-spoiler validation, never quizzing on rising numbers, provably-fair dice arithmetic shown in the ledger) folded back into the generators. **2.3.1** matured Conquest through live founder play (turns resolve the moment every player has committed orders, Risk-scale armies, a dice ledger that shows its arithmetic, match replay), moved the games clock to Pacific so no US player's game day ends before their own midnight, completed a full **accessibility pass** (keyboard focus visibility, dialog semantics, labels throughout), hardened every upload endpoint with magic-byte content sniffing — and gave **every location its own full newsfeed**: the story count on any place's panel opens an infinitely-scrolling, date-filterable river of everything that ever happened there. **2.3.2** put Marco in the war room: a live brief on every Conquest table, a coaching method, a general's playbook and a guided walkthrough for intimidated newcomers; Conquest itself gained rotating initiative, two-sided sabotage, table talk with Marco at the checkpoints, borders that are exactly what the board draws, and month-long Epic wars through changing news weather; the feed gained a publisher blacklist and Google Safe Browsing on every host; and a production ledger review removed a phantom search fee from one AI source's cost accounting. **2.3.3** answered the first week of real players: the map now recovers on its own when a backgrounded mobile browser drops its graphics context (previously it returned unusable until the phone was rotated), the hunt panel moved clear of the mobile controls, a finished hunt updates the Games page immediately, the weekly Bullshit score counts the round that earned it, joining a Conquest table asks first — and in the operator console the query audits finally show Marco's replies, with a sortable, full-width user table beside them.

The crawler (2.4–2.5, September 2026). Version 2.4.0 "Marco reads the news" gave the engine its free second route: Marco the crawler reads the front pages of the publishers already on the map every eight

hours and writes their stories from the article markup itself — the same corroboration as a vendor's find, with no model call and no search fee — with a curator's cockpit (start, pause, the live queue, every rejection and why, the publisher roster with its avoid list and home cities) and an operator's report of how much of the map is now crawled rather than bought; the same release made the traffic console's repeat-visitor figures honest on windows of any length. **Version 2.5.0 "One kind of news"** let the map be asked for one thing: a **topic filter** — the legend's eighteen colors as a picker in the filter bar, as clickable swatches, as the Layers chips — carried through the map's markers, the story strips, every place's story list and the Newsfeed; a **National** tab that trends one country by where each story happened, starting on the reader's own; and a crawler that keeps only what a small model can file under a topic with confidence, reading each story's mood in the same call, with a cleanup that brought its archive up to the same standard. **2.5.1 "Who's looking"** put the site's pulse in the footer — how many stories, threads and places the map holds and how many people are looking at it this hour, each a link — and a **Visitors** layer on the map itself: a tiny dot wherever someone is looking from, city-level and anonymous, from the edge network's own geolocation (both now staff views); the crawler learned each publisher's language and leans its roster toward the readers', reads articles from the first minute of a run instead of after every front page, and the feed engine's spend follows the news down the ladder to the places no publisher covers. **2.5.2 "Fresh to the hour"** gave the crawler two more senses: a **scout** that asks the free news feeds who covers each place and puts those publishers on the roster — a town of a thousand people gets its local paper read without a vendor search — and a **headline lane** that lapped the top forty publishers every hour beside the big run, so the big sites were fresh to the hour while the long tail was walked; the story strips learned what each is for (world and politics on Global, never a school sports score); the Locations directory grew sortable freshness columns and a per-place crawl for curators; map popovers lead with the week's best story instead of the archive's most-corroborated; the operator console gained the corpus against its disk; and in Conquest, drones cost half. **2.5.3 "The day's pace"** made the crawler a steady presence rather than a burst: a daily story budget paces it across the day, its roster leans to the readers' countries, the story strips rank by corroboration first, a wire copy on a local station's site places itself or is not kept, the ten top-ranked storylines carry a fresh AI analysis every day, and the operator console gained the headline

lane's spend, the places crawled and the honest count of people on the map.

Closer to the source (2.6.0, September 2026). Version 2.6.0 "Closer to the source" handed the frontier models' place searches to the crawler: a **location lane** reads, every hour, the places short of their day's news through their own newsrooms; a one-time grounded question finds newsrooms for the places the free scout could not staff; and the scheduled place searches are off — 86% of their spend had bought stories for places the crawler already covered. Every publisher now carries a **reliability rating** drawn from Wikipedia's perennial-sources consensus, with the wires, networks and papers of record leading the crawl and the story strips; a frontier model's new story needs a picture, and a story can borrow a verified picture from a copy that corroborates it. The crawler's topic judge now works in batches across publishers, which made a 10,000-story day affordable, and the headline lane laps eighty publishers every half hour. The video layer checks each channel's live broadcast every half hour, retired five channels and gained 74; search starts at two letters, so "AI" works; the operator console separates people from scripts and says where visitors come from; the welcome page leads with the games; and Conquest gained planes that strike two borders out, hackers and a cyber order, tables that start when their players say so, and fair starts dealt on the board's own borders.

Ask Mappit (2.7.0, September 2026). Version 2.7.0 "Ask Mappit" took the map to where people ask their questions. An **MCP server** at <https://mappit.ai/mcp> lets Claude, ChatGPT, Claude Code, Cursor, VS Code and any other assistant that speaks the Model Context Protocol read Mappit — briefings for any place, country or the world, every story and storyline with Mappit's analysis, the digests, 14.9 million businesses and landmarks, and the live layers — free and without an account, with every story linked to its publisher and to the map. A browser opening the address shows each assistant's setup, Marco explains it on request, and an operator console measures every call, including the questions Mappit could not yet answer. Claude was its first critic, and what it found shipped the same day: the crawler's small judging model now also says **where each story happened** whenever the crawler could only guess — moving a story to the right city, or declining one that happened nowhere in particular — and "this day in history" events are **pinned where they happened** rather than on the nearest city. The corpus chose **quality over quantity**: publishers rated unreliable or deprecated, and those a

curator set aside, are refused at the door, and 5,470 of their stored stories were removed. Anthropic's Claude and xAI's Grok joined Google's and OpenAI's models on the world scan, OpenAI's rejoined the history pass, and the privacy policy now says plainly what the logs keep.

Lay of the Land (2.8.0, October 2026). Version 2.8.0 "Lay of the Land" taught Marco geography. A reference library of open data — 252 countries, every state, province and county, the 235,917 towns and cities of 500 people or more, names in about 30 languages, 1.56 million dated statistics, and land borders with their lengths, drawn from GeoNames, the World Bank, the US Census Bureau, Wikidata, Natural Earth, Unicode CLDR and the CIA World Factbook — now stands behind it, so a question about a place's population, borders or economy gets an answer with the year and source of every figure instead of a larger model's memory. A **Borders** layer draws national and state lines on all six views of the map and on the globe. The crawler's small judging model, which since 2.7.0 says where a guessed story happened, can now also correct the **home towns of publishers** whose stories it keeps placing elsewhere (the first pass is the founder's), and the stories placed before it existed were checked the same way, every change undoable; places that share a name, such as Washington the state and Washington, D.C., are told apart; and a place's panel reads in the reader's language and the place's. Claude Sonnet 5.5 took over Marco's harder questions and the games' writing at its predecessor's price. Conquest changed with its players' feedback: a one-week war's turn waits at most a day for a missing general, an epic waits for everyone, planes need a scientist, a cyberattack takes two steps, one specialist of each kind counts in a province, and a cyber event card arrives with real hacks in the news. And the design for what the library and the borders make possible next — a statistics map — was written (§9).

Almanac (2.9.0, October 2026). Version 2.9.0 "Almanac" put the numbers on the map. Seven **map apps** — World Economy, People, US Economy, Health, Conflict & Crisis, US Elections and News Atlas, 62 views in all — colour countries, US states, counties and congressional districts by 80 indicators from 27 open-data providers, from the World Bank and the US statistical agencies to the Uppsala Conflict Data Program's conflict events, the WHO's outbreak reports and FEMA's disaster declarations, and replay them through time (§4.12). US Elections follows every 2026 Senate, governor and House race with the candidates' fundraising and, for Senate and governor, Mappit's own poll

averages, replays results back to 1976, and is set to read four states' election-night feeds (Georgia's once the state publishes its 2026 results address). A question typed into an app's panel picks the view that answers it; Marco looks the figures up, ranks them, cites them and shows them on the map; and assistants connected through MCP can read them too. Before release the work went through an adversarial code review and two checks of the real app against real data, and every confirmed finding was fixed. In the same week the crawler was retuned and went from about 8,500 to about 15,000 stories a day; Claude Haiku 5.5 took over Marco's everyday answers and Conquest's event cards at a tenth of its predecessor's price; the pages search engines read for a place or a storyline began carrying Mappit's analysis of it; Borders became a default layer; the legend gained each layer's real colour scale; Where's Marco rests the cities it used in the last two weeks; ties on the Where's Marco and Bullshit leaderboards go to the faster player; and Conquest's dispatch names whose province is whose.

Three Ways In (2.9.1, October 2026). Version 2.9.1 "Three Ways In" rebuilt the front door. The welcome page now presents Mappit as what it is, a multi-modal newsfeed and an alternative to feeds ranked by outrage: one world, read on the map, in the Newsfeed and through Threads, with the newest stories and storylines arriving live and a replay from the map apps. Threads list the storylines a reader can read in their language, and a headline's own script now settles its language when its page declares another; a long search runs on its key words, and story searches answer about fourteen times faster; every story an AI model finds carries a picture from the publisher's own page; a shared Threads page or map app view previews as itself; map app replays rescale where one colour scale would paint whole years alike; the map opens on a satellite view with labels; and Conquest draws who moves first.

Cost of Living (2.9.2, October 2026). Version 2.9.2 "Cost of Living" filled the US Economy app with the household budget, every figure from an official source: inflation, from the Bureau of Labor Statistics for each region of the country back to 1967 and from the Bureau of Economic Analysis state by state; the weekly price of a gallon of regular and the monthly price of a kilowatt-hour, from the Energy Information Administration; new homes permitted, from the Census Bureau; mortgage rates and the share of mortgages in foreclosure, through the 2010 peak, from the Federal Housing Finance Agency's National Mortgage Database; and personal bankruptcy filings county by county,

from the federal courts. Directions now come from openrouteservice, with real walking and cycling routes where the old public server had quietly given the car's, and any AI assistant connected to Mappit can ask for them: towns, businesses, landmarks or street addresses, with stops on the way and a link that opens the route on the map. US Elections gained Mappit's average of the House polls, district by district. A Conquest war can run on for more weeks when its players want it to, and the standing table Mappit kept open left the lobby. For visitors arriving from the ads, nearly all on phones, Mappit now offers their own city's news from the welcome page, a free morning digest of it, and a sign-up of an email and a password.

Ledger (2.9.3, October 2026). Version 2.9.3 "Ledger" taught the US Economy app what households owe: debt per person, car loans, credit cards, student loans and mortgages, the share of each that is seriously late, and medical debt in collections, by state and county, from the Federal Reserve Bank of New York and the Urban Institute, in eleven new views. Search engines now get a page for every 2026 race, every state's figure for every statistic and every country's conflict map, the searches a study of Mappit's competitors found them winning with one page per place, and hear about each new storyline analysis and map app view as it appears; a place's page is titled as its local news. The top ten storylines' analyses are rewritten every four hours while new stories arrive, under a daily ceiling that now counts every one. AI agents find the MCP server on their own, through a server card, links in the home page's headers and an API catalog, and it was listed in its first directories; directory test clients that probe it with malformed calls are told what to correct. A map app period now looks the same however it is reached: each is coloured from its own figures alone, a trend line over the timeline shows the change over time, and a replay ends on the map it opened with. Links that are not articles, such as tag pages and publishers' test servers, are refused, and Mappit has one contact address, support@mappit.ai.

Climate (2.9.4, October 2026). Version 2.9.4 "Climate" added an eighth map app, **Environment**, in ten views: temperature against normal and average temperature for every country every month since 1950, from Copernicus's ERA5 record through Our World in Data; carbon dioxide per person, the renewable share of electricity and fine-particle pollution; US temperature and rainfall against the twentieth-century normal by state since 1895 and by county since 2015, from NOAA; the

share of days with unhealthy air in each county with a monitor since 1980, from the EPA; and drought. It opens with the live weather and air-quality layers on. The US Economy app gained the weekly price of diesel from the Energy Information Administration, what farms spend on fertilizer per acre from the USDA's Census of Agriculture (no free source prices fertilizer by state), and job growth from the Bureau of Labor Statistics, by state every month since 1991 and by county every quarter since 2014, each figure naming the jobs added or lost; the World Economy app gained how much fertilizer each country's farms use. That makes 98 views from 43 providers, all of them open data looked up directly, so the release adds nothing to Mappit's AI spend. Signed-out visitors to the welcome page are now offered the free Daily Digest first, with the map a quiet link beside it; and signing out no longer leaves the next person on the device the previous reader's open app, selection and place on the map.

Season Two (2.9.5, October 2026). Version 2.9.5 "Season Two" gave Conquest its second season on the day its first three-week war ended, a war whose table had checked the board several times a day for three weeks. The board now wraps east to west, so no general starts in a corner, and each general's view opens with them in the middle. Every seven turns is a week of the war, and each week prices, recruits, orders and income grow, the cheap things fastest, so a hoarded treasury loses value and a long war turns to its expensive tools. Missiles strike an asset three borders away; satellites are bought for attack or for defense; assassins, the dearest unit in the game, reach any province on the board and need two attempts to kill a scientist, hacker or spy, or one on a day an assassination makes the news. A general can now be told, by a push to a phone or a computer or by email, when the table is waiting only on them, and when a war ends its dispatches open with the war in review: the correspondent's account of the whole war above a table of every general's points, the arithmetic shown. Mappit Plus became what it had been all along, a contribution: Mappit carries no ads, so the welcome and pricing pages now say plainly that readers keep the map running, free members' Daily Digest closes with a chip-in, a short letter goes to them twice a month once an administrator approves it, and Plus members can write a storyline's analysis. The US Elections app gained the early vote, every state's ballots so far against its 2022 turnout from the University of Florida's Election Lab, used with its permission, and North Carolina's by county, with the early voters' party registration;

Marco coaches a Conquest table from the board's own borders, on Sonnet; and every storyline can be shared from its card.

Frontier (2.9.6, October 2026). Version 2.9.6 "Frontier" named what the map apps had become: **frontier geo-intelligence**. The Environment app now draws places as well as colours: every nuclear reactor in the world from Global Energy Monitor's tracker, the daily power level of each US reactor from the Nuclear Regulatory Commission, the 4,874 data centers mapped in OpenStreetMap, and the 95 frontier AI campuses Epoch AI tracks from satellite images and permits, each growing on the map quarter by quarter, alone or over household electricity prices and each country's grid carbon; beside them, how each country makes its electricity, from coal to solar to nuclear. Every site's card says how far it is from the reader, and the map's ruler is open to everyone. Each app's panel gained an **Analysis** tab: a Plus supporter asks, and Claude Sonnet, thinking hard, reads the view's figures, its biggest movers, its trend, the region selected and the news Mappit placed there, and writes what it sees, for every reader to read, at a few cents an analysis. A map app opened from a link outside Mappit now starts without the news pins over its colours, and Conquest's rulebook leads with the game itself. And Conquest no longer needs a second player: a table can reserve a seat for **Marco**, who plays a solo war against you or joins a bigger one as a wildcard, at the level you pick or one that adapts to you, and on your side if you ask.

9. Roadmap and vision

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

- **Deeper and broader coverage** — more places and more languages read through their own newsrooms; the major wires and papers of record that close their pages to crawlers read through the news sitemaps they publish; and more of the physical and economic world layered onto the map (additional markets, additional live feeds).
- **A cinematic map, deepened** — route fly-throughs, the chase camera, and Rewind shipped in 2.2.0; the machinery behind them

(a single camera director owning all choreography) is built for what's next: guided fly-throughs of a *storyline's* geography — the camera touring the cities a developing story spans — and richer Rewind playback over months-long windows as the archive compounds.

- **A statistics map, continued** — the map apps of 2.9.0 are its first part (§4.12): countries, US states, counties and congressional districts shaded by the numbers and by Mappit's own news, played back through time, with Marco ranking regions and showing the answer on the map. Next is the US midterm election on 3 November, when the US Elections app reads the results feeds of four states and shows each one's results once they have been checked. After that, the same design shades the first-level regions of every country, not only the US states, once their shapes are served; and it adds a dossier for any region, a "Silent Map" of the places the world's newsrooms overlook, a headline map, regional editions of the news, and more figures gathered from official pages once the first prove themselves. A licensed source of race calls is under consideration. It is a plan, not a promise: several of its decisions — among them how to draw disputed borders — are still open.
- **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, in every assistant** — since 2.7.0 AI assistants read the corpus directly through Mappit's MCP server, and since 2.9.0 the map apps' statistics too. Since 2.9.3 agents can find the server on their own, and it was listed in its first directories. Next: more directory listings; sign-in, so an assistant can save pins, follow places and ask Marco for a member; the map itself drawn inside the conversation; and a crawler that reads first where people are asking.

The through-line: the map is the beginning, not the product. The product is a permanent, growing, sourced, 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 or other asset**, 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 any token or 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.