

ThreadFox / Preview

The million-view case study and your first session.

AMF Indiana · September 2026

A practical course, publishing toolkit and reseller edition by AMF Indiana. Edition 1.0, September 2026. Purchase price: \$49 USD, one time.

The owner directed the operation. AI handled most posts and comments. That is the operating model this course teaches: give an agent a useful assignment, the right sources, an authorized browser and a clear record of what happened. You should not need to copy every draft between windows or keep reminding the agent which posts it already sent.

This is a complete working curriculum, not a promise of effortless reach. You will build a small publishing workspace, run the included local tools, commission a useful post, authorize an agent to publish it, handle the first comments and review the evidence. The Python toolkit runs locally without an API key. AI services and a browser-capable agent are separate; access, subscriptions and usage limits depend on your provider. A chat-only assistant can prepare the same work, but cannot claim to have clicked Publish.

The course includes a historical case study, editable slides, narrated demonstrations, prompts, Python and JavaScript examples, Windows and Mac/Linux setup launchers, agent instructions, CSV worksheets and a resale package. Demonstrations use clearly identified fictional fixtures. They do not send Reddit posts.

Your first session

Unzip the kit into a new folder. Open START-HERE.html. If you prefer the no-code track, open the prompt library and community worksheet, then use your existing AI chat and browser. For the agent track, run `setup.cmd` on Windows or `sh setup.sh` on Mac/Linux. These launchers check Python and create a local workspace; they do not install a browser extension, purchase a subscription or sign in for you.

Work through the modules in order for your first post. Experienced operators can start with Module 3 and the AI operator contract. Use the answer key to check your work. A successful first session ends with one useful published post and a saved permalink, or a precise blocker. A plausible draft alone is not a successful publication.

01 - Understand what actually worked

The case study starts with a real operational distinction: the human supplied direction and judgement; AI performed much of the execution. The owner's account is that AI handled most posts and comments. The saved system includes publication receipts, post observations and confirmed comment records. The transferable lesson is the division of labour and the feedback loop, not a magic prompt that causes a particular view count.

On September 23, 2026 at 13:03:54 UTC, the saved dashboard snapshot recorded 1,027,710 cumulative views across 140 measured posts. By 23:38:24 UTC, its recorded total was 1,028,774 across 144 measured posts. These are sums of recorded post views. They are not unique readers, verified visits to a website, sales or earnings. Some counts were rounded and some retained their earlier observation time. The set of measured posts changed, and includes posts created before the three-day observation period. Consequently, this record supports 'over one million recorded post views'; it does not establish 'one million new views earned in exactly 72 hours'.

A publication-date audit gives the strongest dated version: September 16-23 contains 139 dated tracked posts, with view observations for 109 totaling 1,009,664. September 20-23 contains 114 posts and 996,022 saved views. September 21-23 contains 84 posts and 349,066 saved views. The last seven calendar days contain 124 posts and 998,921 views. These are approximate cumulative views of each dated cohort, not a fully observed hourly acquisition curve. The eight-day cohort supports the headline: Over one million recorded views across eight days of posts.

The useful mechanics were repeatable: find a question an audience already cares about; provide a concrete answer or worthwhile source; choose a relevant placement; publish; observe actual reactions; improve the next assignment. Community eligibility and subject fit are separate. A large audience is useful only when it is relevant and accepts that type of submission. A successful post in one community does not authorize copying it everywhere.

Build an evidence ladder in your own work. A prepared draft proves preparation. A submit action proves an attempt. A reopened permalink with matching account and text proves the post exists at that observation. Owner visibility does not prove public visibility. A view count proves that the platform displayed that counter, subject to its precision. A confirmed payment proves a transaction. Never collapse these into one success label.

Do this now: write one sentence describing your publication's audience and one sentence defining a useful reader outcome. For example: 'We help first-time container gardeners diagnose common watering problems. A useful post lets someone decide whether to water today.' Then write the business outcome separately: a visit to your guide, a subscribed reader or a purchase. This prevents the agent from optimizing a view counter while forgetting the purpose.

Exercise: your dashboard rises from 10,000 views on ten posts to 40,000 views on twenty posts. Can you report 30,000 newly earned views during that interval? Answer: not without aligning the same post cohort and observation times. Some increase could come from newly imported historical counts. You can accurately report the two snapshots and their coverage.

Deliverable: a one-page brief with audience, reader outcome, business outcome, allowed claims and evidence limitations. Store it in `workspace/brief.md`.

What comes next

The complete kit continues with operator setup, community research, AI writing and browser publication, comments, coding, measurement, removal handling, weekly operations, resale and affiliates.

\$49 USD: <https://threadfox.vip/>