

FRONT PAGE

The blog that *writes itself.*

AutoBlogMatic lives inside WordPress and publishes researched, illustrated, SEO-ready articles on a schedule you set — with your own AI key and no monthly fee.

THIS MONTH

1,842

articles published on autopilot, across customer sites.

COST TO RUN

\$0

per month on a free AI tier. The licence is paid once.

WHY BLOGS STALL

Publishing every week is a full-time job nobody *budgeted* for.

Research, outline, draft, edit, illustrate, optimise, publish. Three to four hours per article, every article, before a single visitor arrives. Most owners stop at week three.

“The tools that help charge monthly — and cap how much you can publish.”

Koala, Journalist.ai and Autoblogging.ai rent words by the month. Stop paying and the engine stops. The blog you spent a year building goes quiet.

THE NUMBERS

3–4h

per post, by hand.

\$150–300

per month, forever, for the tools.

2×

monthly publishing is outranked by daily publishing.

HOW IT RUNS

Install. Add a key. Set the hours. *Close the tab.*

01 Install Upload the plugin, paste the one-time licence key. Live in minutes.	02 Add a key Gemini, OpenAI, Claude, Groq, Mistral, DeepSeek and more. Stored encrypted.	03 Set topics Choose subjects, posts per day, active hours and weekend pause.	04 Publish Researched, illustrated, schema-ready posts go live on schedule.
---	---	--	--

80+ POSTS/DAY CAPACITY	1,800 WORDS PER ARTICLE	10+ AI PROVIDERS	\$0 MONTHLY FEE
--	---	--	---------------------------------------

ONE PAYMENT

Buy it once. Run it *free*.

\$49 at launch, \$99 standard. Single, Business (3 sites) and Agency (unlimited) licences.
Automatic updates for life, 7-day money-back guarantee.

“Your AI key, your quota. Most blogs run on a free Gemini tier.”

No article caps, no renewals, no per-seat pricing. The only recurring cost is one you already have — a free AI key.

ON THE RECORD

4.9★
early-user rating.

1,842
articles published this month.

\$0
outside funding. Bootstrapped.

EVERY ARTICLE NEEDS AN IMAGE

The bottleneck is the *GPU*, not the writing.

TODAY

Featured images come from the user's AI provider or free stock. Generation is capped by their quota, the same images repeat across posts, and internal-link suggestions need an embedding pass over the archive that has nowhere to run.

WHAT WE WANT ON GPU

- One article-specific featured image per post.
- Embeddings for topic clusters and internal links.
- Batch 30 posts per site, 10–50 sites concurrently.
- Weights restored from a snapshot, then scale to zero.

1

IMAGE PER ARTICLE

30–120s

PER SITE BATCH

10–50


CONCURRENT SITES

0

IDLE GPU HOURS

FROM TOPIC TO PUBLISHED

Topic → draft → image → schema → *live.*



GPU SANDBOX

Image and embedding models with weights loaded from a snapshot, so a 30-post batch starts in seconds instead of downloading models per run.

CPU SANDBOX

WordPress publishing, schema rendering, link checks and the scheduler. Bursty per site, so per-second billing fits.

WHAT WE NEED

\$10,000

in Daytona credits to run image generation and embeddings for every post, on GPU sandboxes.

WHAT WE WILL RUN

- GPU sandboxes for featured images and topic embeddings.
- 30-post batches per site, 10–50 sites concurrently.
- Benchmarks: GPU image generation vs provider API cost.

WHAT YOU GET BACK

Cost and latency benchmarks for per-post image generation, real usage data from production sites, a written case study and developer content about running WordPress AI publishing on Daytona.