

LIGHTS OUT

Research & Production Notes

BY MARCELO KANHAN — APRIL 2026

This document accompanies the featured article *Lights Out* on kanhan.com.br. It documents the sources, editorial decisions, and production methodology behind the piece. The article was produced using AI-native editorial infrastructure — Claude Code as autonomous production system, with the author providing editorial direction, quality gates, and final judgment at every stage.

PRODUCTION METHODOLOGY

The article was produced through a pipeline where the author provides seeds, editorial direction, and quality checkpoints, and Claude Code handles research, drafting, source verification, and production. The process is documented because transparency about how things are made is a core principle of this practice.

Pipeline: Markdown files with frontmatter state in a local directory. A nightly cron job advances eligible pieces. The author reviews at checkpoints. No database, no framework, no UI — the LLM reads, understands, acts.

Iterations: 3 major drafts.

Fact-check: Medvi case independently verified against NYT, Morning Brew, Health Data Consortium, Futurism, FDA records. Claims qualified per epistemic boundaries (verified vs. reported vs. alleged).

SOURCES BY SECTION

I. THE THRESHOLD

- **SWE-bench Verified leaderboard (March 2026):** Claude Opus 4.5 at 80.9%, Gemini 3.1 Pro at 80.6%, GPT-5.2 at 80.0%. Top six models within 1.3 points.

Source: vals.ai/benchmarks/swebench; llm-stats.com

- **SWE-bench Verified retirement:** OpenAI flagged training data contamination across all frontier models, February 23, 2026. SWE-bench Pro replaced it with scores in the 55-58% range.
Source: openai.com
- **Claude Mythos Preview (April 7, 2026):** 93.9% SWE-bench Verified, 77.8% SWE-bench Pro. Not commercially released — deployed exclusively through Project Glasswing for defensive cybersecurity. Discovered thousands of zero-day vulnerabilities including a 27-year-old OpenBSD bug.
Source: anthropic.com/glasswing; red.anthropic.com/2026/mythos-preview; fortune.com (April 7, 2026)
- **Anthropic on Claude Opus 4.5:** “the world’s best coding model, leading on SWE-bench (72.5%) and Terminal-bench (43.2%)” [May 2025 announcement].
Source: anthropic.com/news/claude-4

II. THE AGENTIC LOOP

- **StrongDM Software Factory:** Three engineers, ~32,000 lines (16K Rust + 9.5K Go + 6.7K TypeScript), since July 2025. Charter: “Code must not be written by humans. Code must not be reviewed by humans.”
Source: strongdm.com/blog/the-strongdm-software-factory (Feb 6, 2026); simonwillison.net/2026/Feb/7/software-factory
- **Scenarios as holdout tests:** “We repurposed the word scenario to represent an end-to-end ‘user story’, often stored outside the codebase (similar to a ‘holdout’ set in model training), which could be intuitively understood and flexibly validated by an LLM.”
— StrongDM blog
- **Digital Twin Universe:** “Behavioral clones of the third-party services our software depends on — replicating their APIs, edge cases, and observable behaviors.”
— StrongDM blog
- **Token spending metric:** “If you haven’t spent at least \$1,000 on tokens today per human engineer, your software factory has room for improvement.”
— Justin McCarthy, CTO, StrongDM
- **Stanford Law:** “When software is ‘grown’ rather than written, when replicas stand in for real services, and when quality is measured by probability rather than certainty, who is responsible?”
— Eran Kahana, CodeX, Stanford Law School (Feb 8, 2026). Source: law.stanford.edu
- **Cursor 3 (April 2, 2026):** “a bet that you’ll manage agents, not write code” / “We will continue to invest in the IDE until codebases are self-driving.”
Source: cursor.com/blog/cursor-3; the-decoder.com (April 3, 2026)

III. SOFTWARE WITHOUT VERSIONS

- **CI/CD context:** Continuous Integration / Continuous Delivery — the practice of merging, testing, and deploying code automatically and frequently.
- **Cursor self-driving metaphor:** "We will continue to invest in the IDE until codebases are self-driving."

Source: cursor.com/blog/cursor-3

IV. ATOMIC APPLICATIONS

- **Ephemeral UI:** AI-generated interfaces created in context, disappear when done.
- **Andreas Kirsch, "The Flawed Ephemeral Software Hypothesis":** Making code generation cheaper shifts the bottleneck to validation, integration, and ergonomics. Software isn't ephemeral; it's malleable.

Source: [iSolutions/Medium](https://isolutions.medium.com)

Source: blackhc.net

- **75% projection:** 75% of new enterprise applications AI-generated or zero-code by 2026.

Source: codeconductor.ai

V. THE HUMAN AS SEED

- **FANUC lights-out factory:** Yamanashi, Japan. Robots building robots since 2001. 50 units per 24-hour shift, unsupervised up to 30 days.

Source: [Wikipedia](https://en.wikipedia.org/wiki/FANUC); [Assembly Magazine](https://www.assemblymagazine.com) 2019

- **Addy Osmani:** "Pour your mentorship into the spec" / "if the agent produces something that technically meets the spec but doesn't feel right, trust your judgement."

Source: addy.substack.com (Jan 19, 2026)

- **Thoughtworks:** "Not yet a systematic way to evaluate specs."

Source: thoughtworks.medium.com (Dec 10, 2025)

- **GitHub:** "Your role isn't just to steer — it's to verify."

Source: github.blog (Sep 2, 2025)

- **Content marketing automation:** "Some teams configure fully automated workflows where content moves from topic identification to publication without human intervention." "Slop" — Merriam-Webster's Word of the Year 2025.

Source: trysight.ai; [Jasper](https://jasper.ai), 2026 State of AI in Marketing

VI. THE DRIVE BENEATH THE AUTOMATION

- **Medvi / Matthew Gallagher:** NYT profile by Erin Griffith (April 2, 2026). \$401M revenue in 2025 (NYT reviewed financial data). \$1.8B is revenue run-rate

projection, not valuation. Two employees. FDA warning letter February 20, 2026. Fabricated physician profiles confirmed by multiple outlets. Gallagher acknowledged AI-generated photos to the NYT.

Sources: nytimes.com; morningbrew.com; healthdataconsortium.org; futurism.com; FDA

- **Dark factory frameworks:** BCG Platinion, HackerNoon, darkfactory.dev.

Sources: bcgplatinion.com; hackernoon.com; darkfactory.dev

- **Lights-out manufacturing history:** IBM keyboard plant (Texas, 1980); FANUC (Yamanashi, 2001).

Source: Wikipedia; Assembly Magazine 2019; IMTS 2026

EDITORIAL DECISIONS

- **Medvi reframed as cautionary tale.** Original draft presented Medvi as an example of successful leverage. Fact-check revealed fabricated physicians, FDA action, and data breach. Reframed to illustrate what happens when the factory runs without judgment.
- **Lights on/off in every section.** The toggle concept — readers choosing between the before and after — is a structural argument, not a gimmick. Each section presents the same scenario from both sides.
- **Section VII added.** Original scope was six sections. The closing (“The Lit Room”) was added to synthesize the full arc.
- **Technical jargon reduced.** v3 pass softened language for non-technical readers (CI/CD explained, benchmark context simplified).
- **Claude Mythos added.** Anthropic’s April 7, 2026 release incorporated as the most recent data point on the capability trajectory.

This article was produced using Claude Code as autonomous editorial infrastructure. The author provides seeds, editorial direction, and quality gates. The production pipeline — research, drafting, source verification, fact-checking, and formatting — operates with high autonomy between checkpoints.