




Welcome to the **July 2026** edition of **The AI Advantage**.

This month: SpaceX acquires Cursor for \$60B, OpenAI and Broadcom unveil their Jalapeño AI chip, Claude Opus 4.8 launches for the agentic era, Microsoft drops the MAI model suite -- and four must-read posts from the blog.



THIS MONTH ON THE BLOG

Four Posts You Can't Miss in July

Andrew Murphy · July 2026 · Blog Preview

July on the blog is about results, agency, and income systems -- four posts that connect the tactical (how to automate, how to create content faster) with the strategic (how AI compounds your affiliate income over time). Here's what's live:

Week 27 -- July 6

AI Automation ROI: The Real Numbers Small Businesses Are Seeing in 2026

The average AI automation ROI for small businesses in 2026 is 280% to 520% in the first year for focused deployments. This post breaks down where the returns come from, the tools generating the best results, and the five-step framework for getting your first automation live this week.

Week 28 -- July 13

Agentic AI in 2026: How Autonomous Agents Are Replacing Entire Workflows

AI agents plan, act, and complete multi-step tasks without waiting for instructions. 40% of business applications will have agents built in by end of 2026. This post covers what agents can actually do today -- and the no-code tools making it accessible without a developer.

Week 29 -- July 20

One Afternoon, One Month of Content: The AI Social Media Blueprint

Short-form video delivers 49% higher conversion rates than text. This post is the exact system -- five hours, one blog post, one recording, and the right AI tools -- that produces 40 pieces of content scheduled across six platforms for a full month.

Week 30 -- July 27

The 2026 Affiliate Marketer's AI Playbook: How Top Earners Scale With AI

96% of top affiliate creators now use AI in their workflow. This post is the exact playbook they're using -- recurring commission programs, content volume, email sequences, and the realistic income timeline from month one through month twelve.

All four posts go live on their scheduled dates at andrewmurphy.online/blog-home -- with affiliate links so you can try every tool directly.

[Visit the Blog ->](#)



ACQUISITION

SpaceX Buys Cursor for \$60 Billion -- The Biggest VC-Backed Deal in History

AI Industry · June 16, 2026

SpaceX announced it is acquiring Cursor -- the AI coding assistant built by San Francisco startup Anyosphere -- for \$60 billion in an all-stock deal. It is the largest VC-backed startup acquisition in history, closing in Q3 2026.


Why Cursor? By mid-2026, Cursor was generating approximately \$2.6 billion in annualized B2B revenue. The platform had become the dominant AI coding environment for professional developers worldwide -- a product category with extremely high retention and expanding enterprise contracts.

The strategic picture: SpaceX merged with xAI earlier in 2026, combining Elon Musk's AI company with its compute infrastructure. The Cursor acquisition gives the combined entity a direct product in the hands of millions of developers daily -- a distribution channel no other AI lab has matched.

"\$60 billion for a coding assistant is not about the tool -- it is about daily developer habits and the distribution that comes with owning them."

What it means for affiliate marketers in the AI tools space:

AI coding tools are the highest-growth, highest-retention category in software. The SpaceX/Cursor deal validates the category at a scale that will accelerate enterprise adoption. If you are promoting AI tools to a business or developer audience, coding assistants -- and the productivity narrative behind them -- just got a massive credibility boost. Content on AI coding tools will rank and convert better than ever.



HARDWARE

OpenAI and Broadcom Unveil Jalapeño -- A Custom AI Chip Built in Nine Months

AI Infrastructure · June 24, 2026

OpenAI and Broadcom revealed Jalapeño -- OpenAI's first custom-designed AI inference chip, built from initial design to manufacturing tape-out in just nine months. The companies say it targets 50% cheaper inference costs than current NVIDIA GPU deployments.




Why this matters: Inference cost is the single biggest constraint on how frequently AI tools can run complex tasks. At 50% lower cost per inference, the economics of running AI agents, generating content at scale, and building automated workflows shift meaningfully in favor of smaller operations.

The nine-month development cycle is itself significant. OpenAI used its own models to accelerate parts of the chip design and optimization process -- a recursive loop where AI builds the hardware that runs AI faster. Initial deployment is targeted for end of 2026 in gigawatt-scale data centers.

"Cheaper inference means more AI calls per dollar -- which means more agents, more automations, and more content output from the same monthly budget."

The practical impact for creators and marketers:

Every AI tool you use -- Claude, ChatGPT, Jasper, Make AI steps -- runs on inference compute. As custom chips like Jalapeño come online, the cost of AI tool subscriptions will continue to fall while capabilities rise. The tools you promote today will become more affordable and more powerful through 2027 -- a tailwind for anyone building an audience in the AI tools affiliate space.



NEW MODEL

Claude Opus 4.8 Launches -- Built for the Agentic Era

Model Release · May 28, 2026 · Anthropic

Anthropic released Claude Opus 4.8, the model now powering this newsletter's research and writing workflow. The headline feature is Adaptive Thinking -- the model now decides when to reason deeply and when to respond directly, making it 2.5x faster on simple tasks while maintaining full reasoning depth on complex ones.

Dynamic Workflows in Claude Code allow the model to tackle very large-scale problems autonomously -- multi-step coding and content projects that previously required manual checkpoints can now run uninterrupted. For entrepreneurs using Claude for content production, this means longer, more coherent outputs with fewer interruptions.

Effort control is available on claude.ai: you can now set how much reasoning effort Claude applies to a task. Fast mode runs at 2.5x speed and is 3x cheaper -- useful for high-volume content generation where speed matters more than maximum depth. System Message Flexibility allows updated instructions mid-conversation without losing prompt cache hits.



"Adaptive thinking means Claude now matches its compute to the complexity of the task. Simple requests are instant. Complex agentic tasks get full reasoning depth. You pay for what the task actually needs."

How to use it right now:

If you are running Claude for blog post drafting, use Fast mode for first drafts and standard mode for final editing and SEO optimization. The speed difference is significant -- and the cost saving on volume production compounds quickly. Claude Opus 4.8 is available at claude.ai and via the Anthropic API at the same price as Opus 4.7.



AI MODELS

Microsoft Launches Seven In-House AI Models Under the MAI Brand

Enterprise AI · June 2, 2026 · Microsoft

Microsoft unveiled a suite of seven in-house AI models under the MAI designation -- its most significant move toward AI independence from OpenAI since the two companies began their partnership. The flagship models are MAI-Thinking-1, a reasoning model designed to match premium logical output at competitive token cost, and MAI-Code-1-Flash, Microsoft's first

code-generation model.

MAI-Code-1-Flash generates source code for applications and websites from written descriptions -- no programming background required. For non-technical entrepreneurs who have been relying on developers or no-code tools, this is the most practically significant release of the batch. Early testers are using it to build simple automation scripts, landing page logic, and API integrations without code.

The broader strategy: Microsoft has been investing heavily in AI infrastructure independently of its OpenAI partnership. The MAI suite is designed for Azure deployment at competitive pricing -- positioning Microsoft's cloud as the enterprise AI platform of choice regardless of which foundation model a business prefers.

"The MAI suite signals that Microsoft is no longer betting exclusively on OpenAI. Enterprise AI is becoming a multi-model discipline -- and Microsoft wants to be the platform that runs all of them."

Content opportunity for AI marketers:

The MAI suite launch creates a natural content moment: comparison posts, tutorials, and 'which Microsoft AI model should you use' guides are high-search, low-competition right now. If you produce content in the AI tools niche, write a comparison of MAI-Code-1-Flash vs GitHub Copilot vs Claude for code generation this month. The search volume is building and there are very few authoritative guides yet.





REGULATION

Colorado Rewrites Its Landmark AI Law -- New Version Takes Effect January 2027

Policy & Compliance · May 2026 · Colorado SB 26-189

Colorado's original AI Act -- the first comprehensive state AI law in the US, passed in 2024 -- has been substantially rewritten. Governor Polis signed SB 26-189 in May 2026, delaying the effective date from June 30, 2026 to January 1, 2027, while significantly narrowing the law's scope and requirements.

What changed: The original law's broad risk-based framework -- which required risk management programs, bias mitigation, and impact assessments for any 'high-risk AI system' -- has been replaced with a narrower regime focused on 'automated decision-making technology' that materially influences consequential decisions. Most small business and content creator AI uses fall outside the new definition entirely.

The regulatory trend remains clear: Even as Colorado scaled back, the direction of travel across all jurisdictions is toward disclosure requirements. The White House's National Policy Framework pushes for federal preemption -- one set of rules across all 50 states. The patchwork

continues, but transparency requirements are the common thread in every version of AI regulation being proposed.

"Every version of AI regulation -- federal, state, EU -- converges on the same requirement: disclose when and how you use AI. That disclosure is already good business practice regardless of what the law requires."

Simple action for online business owners: Add a short AI disclosure line to your content footer, email signature, and About page -- something like 'This content was produced with AI assistance.' Platforms like YouTube already require it for AI-generated video. The FTC increasingly expects it for sponsored and AI-assisted content. Getting ahead of compliance now costs 30 seconds and protects you against requirements that are coming whether Colorado's specific law applies to you or not.

Affiliate angle:

Regulation content converts extremely well with business owner audiences. A post titled 'What the Colorado AI Law Means for Your Online Business' or 'AI Disclosure Requirements: What Every Affiliate Marketer Needs to Know' ranks quickly, builds trust, and naturally leads into tool recommendations for compliance-friendly AI content workflows.

