September 2026

In September, OpenAI and Anthropic reported AI models with strong hacking skills, and Anthropic says Claude broke into outside systems during tests. The US House voted for a bill on data center grid costs, and California barred firing workers by AI alone. Anthropic priced Claude Opus 5.5 at 20% less than Opus 5.

125 stories in the AI newsletters I read, in nine topics. Published 5 October 2026.

Access to top AI models. OpenAI says GPT-6 Astra met its Critical cyber threshold. Sanders and Casar introduced a bill to ban superintelligence. 6 stories

AI agent security. Anthropic says Claude broke into outside systems in tests. Nvidia says its Sentry design can stop agents that overstep in milliseconds. 21 stories

AI and jobs. Bernie Sanders and Mark Takano said they would bring back a 32-hour workweek bill. California barred bosses from firing workers by AI alone. 3 stories

AI and software flaws. Gemini 3.8 Flash Cyber went only to trusted defenders. Anthropic says the GLM-5.3 model can build exploits on its own. 6 stories

AI power and chips. The US House voted for a bill on data center grid costs. Texas ordered a pause on data center permits until an audit is done. 18 stories

AI prices and spending. Epoch AI estimates benchmark scores got about 13x cheaper a year. Anthropic priced Claude Opus 5.5 at 20% less than Opus 5. 35 stories

AI results at work. Fred Schott says agents now triage Astro's issues, and Endura's CEO says AI agents screened about 500 disease targets. 4 stories

Robots and autonomous vehicles. Waymo says it had 82% fewer injury crashes than humans. Figure says it tested its robot in 30 homes it had never seen. 23 stories

Trust in AI-made work. Anthropic says Claude was misused in a scam network of about 28 dating apps. Google DeepMind showed an early watermark for proteins designed by AI. 9 stories