

Grounding pages, default training, and the curl zero-day gap

2026-09-02 / 00:06:11

“The sites are optimized for model retrieval, not human readers. Their 'Facts & Grounding Page' titles are addressed to the software that reads them.”

— Seln Oriax, today's narration

Today we're looking at three stories about visibility — what's visible when models make content, what companies do with your data by default, and how many bugs frontier AI systems miss that others catch.

Claude rolled out a content checker that uses C2PA credentials to verify whether Claude helped create or edit a file. Meanwhile, the Claude system prompt (now “Fable 5.1”) got tighter on song lyrics and copyrighted characters — and Simon Willison built an Atom feed to track those changes.

Mistral changed its default training policy for non-enterprise users. And Trellner's research found sites publishing hundreds of thousands of pages designed specifically to be cited by AI models, not read by humans.

Finally: AISLE found six curl CVEs after both Anthropic and OpenAI returned zero in the same codebase.

00:00:04 Content verification and policy tracking

00:01:38 Mistral's default-on training policy

00:02:32 Pages built for models, not people

00:04:18 Six CVEs, zero from the labs
