Gavin Baker's "Land Man" primer for the AI buildout

Recommended Reading from Tech Founders

2026-08-15

Gavin Baker's "Land Man" primer for the AI buildout

By Recommended Reading from Tech Founders • August 15, 2026

The strongest organic recommendation is a TV series used to understand the physical work behind data centers; three additional leads cover orbital-compute assumptions, self-improving research, and collateral risk.

The standout: the physical layer of AI infrastructure

- **Title:** *Land Man* (the title as spoken in the discussion)
- **Content type:** TV series
- **Author/creator:** Not stated in the source
- **Link/URL:** No direct link supplied
- **Recommended by:** Gavin Baker. The episode introduces Baker, asks him to address the energy and data-center supply needed to scale AI, and he answers that the best way to understand the people building data centers is to watch the series. [1]
- **Key takeaway:** Baker's point is "atoms, not bits": the work involves the oil patch, 110° heat, and coordinating thousands of people in remote locations. [1]
- **Why it matters:** This is the most concrete recommendation for understanding AI infrastructure as an operations problem rather than only a model or chip story. [1]

A model for the terrestrial-to-orbital constraint

- **Title:** *AI Compute "Keystone" Model* and its linked public brief
- **Content type:** X post and research brief
- **Author/creator:** Vlad Saigau; the post is associated with Mach33 Financial Group. [2]

- **Link/URL:** X post · public brief [2]
- **Recommended by:** This is a recommendation chain rather than a single direct endorsement: Elon Musk linked Owen Lewis's post while arguing that orbital compute may be the only way to scale AI around 2029; Lewis called the underlying orbital-compute point "literally the most important part" and linked Saigau's work. [3, 4, 5]
- **Key takeaway:** Saigau's model says orbital data-center compute could become necessary for AI scaling in the 2030s. Its own projections put terrestrial capacity at a 111 GW peak in 2031 and roughly 438 GW of a 488 GW 2040 fleet in orbit, while the author stresses that no one—including the authors—has a narrow projection and that the framework is built to evolve. [2]
- **Why it matters:** Read this as an assumption-rich infrastructure framework, not as a forecast to accept wholesale. Its value is that it makes terrestrial capacity, demand, price deflation, and uncertainty explicit. [2]

A concrete pattern for self-improving research

- **Title:** Copacetic's "auto research" X post (no formal title supplied)
- **Content type:** X post
- **Author/creator:** Copacetic
- **Link/URL:** No URL supplied in the source
- **Recommended by:** Tom Blomfield, who calls it a "great tweet." [6]
- **Key takeaway:** The post applies an iterative loop to machine learning: have AI generate research ideas, test them overnight, and hill-climb toward a measurable outcome—discarding moves downhill and retaining moves uphill. [6]
- **Why it matters:** It is a compact design pattern for self-improving ML, where evaluation and repeated experimentation are part of the system rather than a one-off prompt. [6]

A collateral-risk reading for compute finance

- **Title:** John Geanakoplos's academic literature on collateral (the specific paper or title is not named)
- **Content type:** Academic literature / research
- **Author/creator:** John Geanakoplos
- **Link/URL:** No link supplied in the source
- **Recommended by:** Josh Wolfe, one of the panelists identified in the discussion. In the conversation about securitizing AI compute, the transcript directs readers to Geanakoplos's literature. [7]
- **Key takeaway:** The housing crisis is framed as a collateral problem—not merely a bubble or a debt problem—with collateral becoming worthless and the debt following it. [7]
- **Why it matters:** The recommendation is a useful historical check on AI-compute financing: the same discussion flags unknown GPU deprecia-

tion and duration mismatch, so Geanakoplos offers a way to pressure-test whether compute can support long-lived financing assumptions. [7]

Sources

1. Anthropic's $2T IPO, Zuck's AI Manifesto, Nvidia's $500B AI Bet, Grok's Comeback
2. Vlad Saigau (@VladSaigau) on X
3. X post by @elonmusk
4. Owen Lewis (@is_OwenLewis) on X
5. X post by @is_OwenLewis
6. Building And Structuring An AI Native Company
7. Special Edition MOL with Josh Wolfe, Rachel Holt, Scott Belsky, Scott Stanford, and Peter Deng.