

FRED TALKS

6th April 2026

CONTENTS

The AI Great Leap Forward

HAN, NOT SOLO · 2188 WORDS

Brain Food: Credibility is Expensive

SHANE PARRISH (FS) · 617 WORDS

The Axios supply chain attack used individually targeted social engineering

SIMON WILLISON'S WEBLOG: ENTRIES · 357 WORDS

Issue #704

POINTER · 795 WORDS


The AI Great Leap Forward

HAN, NOT SOLO · 05 APR 2026 · [SOURCE](#)

In 1958, Mao ordered every village in China to produce steel. Farmers melted down their cooking pots in backyard furnaces and reported spectacular numbers. The steel was useless. The crops rotted. Thirty million people starved.

In 2026, every other company is having top down mandate on AI transformation.

Same energy.



Backyard Furnaces

The rallying cry of the Great Leap Forward was — surpass England, catch up to America. Every province, every village, every household was expected to close the gap with industrialized Western nations by sheer force of will. Peasants who had never seen a factory were handed quotas for steel production. If enough people smelt enough iron, China becomes an industrial power overnight. Expertise was irrelevant. Conviction was sufficient.

The mandate today is identical, just swap the nouns. Every company, every function, every individual contributor is expected to close the AI gap. Ship AI features. Build agents. Automate workflows. That nobody on the team has ever trained a model, designed an evaluation system, or debugged a retrieval system is beside the point. Conviction is sufficient.

So everyone builds. PMs build AI dashboards. Marketing builds AI content generators. Sales ops builds AI lead scorers. Software engineers are building AI and data solutions that look pixel-perfect and function terribly. The UI is clean. The API is RESTful. The architecture diagram is beautiful. The outputs are wrong. Nobody checks because nobody on the team knows what correct outputs look like. They've never looked at the data. They've never computed a baseline.



Backyard Furnaces

Entire departments are stitching together n8n workflows and calling it AI — dozens of automated chains firing prompts into models, zero evaluation on any of them. These tools are merchants of complexity: they sell visual simplicity while generating spaghetti underneath. A drag-and-drop canvas makes it trivially easy to chain ten LLM calls together and impossibly hard to debug why the eighth one hallucinates on Tuesdays. The people building these workflows have never designed an evaluation pipeline, never measured model drift, never A/B tested a prompt. They don't need to — the canvas looks clean, the arrows point forward, the green checkmarks fire. The complexity isn't avoided. It's hidden behind a GUI where nobody with ML expertise will ever look.

The backyard steel of 1958 looked like steel. It was not steel. Today's backyard AI looks like AI. It is not AI. A TypeScript workflow with hardcoded if-else branches is not an agent. A prompt template behind a REST endpoint is not a model. Calling these things AI is like calling pig iron from a backyard furnace high-grade steel. It satisfies the reporting requirement. It fails every real-world test.

But the most dangerous furnace is the one that produces something *functional*. Teams are building demoware — pretty interfaces, working endpoints, impressive walkthroughs — with zero validation underneath. Some are in-housing SaaS products by vibe coding some frontend with coding agents: it runs, it has a dashboard, it cost a fraction of the vendor. Klarna announced in 2024 that it would replace Salesforce and other SaaS providers with internal AI-built solutions. What these replacements don't have is data infrastructure, error handling, monitoring, on-call support, security patching, or anyone who will maintain them after the builder gets promoted and moves on.

These apps will win awards at the next all-hands. In two years they'll be unmaintainable tech debt some poor soul inherits and rewrites from scratch. The furnace produced pig iron. Someone stamped "steel" on it. Now it's load-bearing.

Meanwhile, the actual product that customers pay for rots in the field. But hey, . The AI adoption dashboard is green.

Reporting Grain Production to the Central Committee

During the Great Leap Forward, provinces competed to report the most spectacular grain yields. Hubei reported 10,000 jin per mu. Guangdong said 50,000. Some counties claimed over 100,000 — physically impossible numbers, rice plants supposedly so dense that children could stand on top of them. Officials staged photographs. Everyone knew the numbers were fake. Everyone reported them anyway, because the alternative was being labeled a saboteur. The central government, delighted by the bounty, increased grain requisitions based on the reported yields. Farmers starved eating the difference between the real number and the fantasy.

You've seen this meeting.

One team reports their AI copilot “reduced development time by 40%.” The next team, not to be outdone, reports 60%. A third claims their AI agent “automated 80% of analyst workflows.” Nobody asks how these were measured. Nobody checks the methodology. Nobody points out that the team claiming 80% automation still has the same headcount doing the same work. The numbers go into a slide deck. The slide deck goes to the board. The board is delighted. The board increases investment.



Reporting Grain Production to the Central Committee

Then someone — there's always someone — builds a leaderboard tracking how many prompts you wrote this week, how much of your code is AI-generated, your ranking versus your team, versus your org, versus the entire company. One day your company announces: stop everything, it's AI Week. Build something with AI. Show what you've got. You think you're done after the hackathon? No no no. Now you have to *promote* it. Daily posts: look what I built, here's how many agents I used, here's how many skills I shipped. Pull in teammates. Pull in strangers. Ask for feedback. “Humbly.”

Your AI usage is now a KPI. You are being evaluated on how much grain you reported, not how much grain you grew. This is Goodhart's Law at organizational scale: when a measure becomes a target, it ceases to be a good measure. The metric was supposed to track whether AI is making the company better. Instead, the entire company is now optimizing to make the metric look better. The beatings will continue until adoption improves.

Killing the Sparrows

The Great Leap Forward's most tragicomic chapter was the **Eliminate Four Pests Campaign**. Mao declared sparrows an enemy of the state — they ate grain seeds, so killing them would increase harvests. The entire country mobilized. Citizens banged pots and pans to keep sparrows airborne until they dropped dead from exhaustion. Children climbed trees to smash nests. Villages competed for the highest kill count. It worked. They nearly eradicated sparrows.

Then the locusts came.

Sparrows ate locusts. Without sparrows, locust populations exploded. The swarms devoured far more grain than the sparrows ever did. The campaign to save the harvest destroyed it. Mao quietly replaced sparrows with bedbugs on the official pest list and never spoke of it again.

Every AI Great Leap Forward has its sparrow campaign.

Middle managers are the sparrows. They're declared pests — too many layers, too slow, too expensive. Flatten the org! Move faster! Let AI handle coordination! So companies eliminate M1s, turn managers into tech leads running pods, and let the teams self-organize with AI tools.



Killing the Sparrows

Then the locusts come. Those middle managers held institutional knowledge — which customer had the weird integration, why the data model had that inexplicable column, the undocumented business rule that kept compliance from flagging every third transaction. That context lived in their heads. Now they're gone, and the AI system they were replaced with needs exactly that context to function.

QA is a sparrow too. "AI writes the tests now." So you cut QA. The AI writes tests that validate its own assumptions — a machine checking its own homework. Senior engineers who mentored juniors? Sparrows. Documentation writers? Sparrows. The ops team that knew how to restart the weird legacy service at 2 AM? Definitely sparrows.

Each elimination looks rational in isolation. The second-order effects arrive six months later, and by then nobody connects the locust swarm to the dead sparrows.

Let a Hundred Skills Bloom

In 1956, Mao launched the (Hundred Flowers Campaign): "Let a hundred flowers bloom, let a hundred schools of thought contend." Speak freely. Share your honest criticisms. The Party wants to hear your real thoughts.

Intellectuals took the bait. They spoke openly.

Then came the (Anti-Rightist Campaign). Everyone who had spoken honestly was identified, labeled, and purged. The Hundred Flowers was a trap — an efficient mechanism for surfacing exactly who knew what, then eliminating them. The lesson every survivor internalized: *never honestly reveal what you know, because it will be used against you.*

Now Meta and a growing list of companies have launched their own Hundred Flowers. The mandate: every employee must build "agent skills" — distill your subject matter expertise into structured prompts and workflows that AI agents can execute. Or even worse, build "agents" using some drag and drop legacy tech that never worked and had already been given up by the leading edge labs back in 2024. Encode your judgment. Document your decision-making. Make yourself legible to the machine.

The stated goal is **distilling** your subject matter expertise. Turn the expert's craft into the organization's asset. What leadership actually wants is to convert individual human capital into organizational capital that survives any single employee's departure.

Employees see the game immediately. If I distill my ten years of domain expertise into a skill that any junior can invoke with a prompt, I have just automated my own replacement. The knowledge that makes me the critical node — the person they call at 2 AM, the one who knows *why* the model does that weird thing for Brazilian entities — is my moat. You're asking me to drain it.

So they adapt to build anti-distillation agent skills, just as the intellectuals adapted after the Anti-Rightist trap.

We are already seeing agent skills built specifically for job security. The *performative skill* looks comprehensive and demos well but omits the 20% of edge-case knowledge that makes it work in production — you are now *more* indispensable, not less. The *poison pill* encodes expertise faithfully but with subtle dependencies on context only you hold — internal wikis you maintain, terminology you coined, data pipelines you own — so removing you causes outputs to drift quietly until someone says “we need to bring them back on this.” The *complexity moat* makes the skill so architecturally entangled with your other work that extracting your knowledge is harder than keeping you around. You are now a load-bearing wall disguised as a decoration.

The campaign designed to reduce organizational dependence on individual experts has now created experts who are *strategically indispensable* — not because of what they know, but because of how they've booby-trapped the system to need them. The flowers bloomed. They're full of thorns.

Meanwhile, the “everyone builds with AI” mandate has turned into a hunger game of scope creep. Engineers use AI to generate designs and ship prototypes without waiting for the design team. PMs use AI to write code and spin up dashboards without filing engineering tickets. Designers use AI to build product specs and run user research without looping in product. Everyone is expanding into everyone else's territory — not because they're better at it, but because AI makes it *possible* and the mandate makes it *rewarded*. The org chart says collaboration; the incentive structure says land grab. What looks like productivity gains is actually a war of all against all, where every function is simultaneously trying to prove it can absorb the others before the others absorb it.



Engineering, PM, and Design scope creep

The Famine Comes Later

The Great Leap Forward's famine didn't arrive immediately. For a while, the numbers looked spectacular. Every province reported record harvests. Leadership was pleased. The requisitions increased.

The famine came when the real grain ran out but the reported grain kept flowing upward.

We're still in the reporting phase. The dashboards are green. Adoption is up and to the right. Every team reports productivity gains that, if summed across the company, would imply engineers are shipping at 300% efficiency while somehow still missing the same deadlines.

Underneath the metrics, it's a race to the bottom. One person builds a skill, so someone else builds a better one. One person demos a prototype, so someone else benchmarks it. Everyone competing to prove, more thoroughly than the next person, that their own role is replaceable. All accelerating. All sinking.

The sparrows are dead. The locusts haven't arrived yet. The flowers bloomed full of poison pills. The furnaces produced pig iron stamped as steel that's now load-bearing. The grain numbers look fantastic.

But it's fine. We're surpassing and catching up.

Oh, and Klarna? The company that loudly announced it would replace Salesforce with internal AI solutions? They quietly replaced Salesforce with another SaaS vendor instead. The backyard furnace couldn't produce real steel. They bought it from a different mill.

The question nobody's asking: what did any of this actually produce?

The answer, when it arrives, will be awkward.

References

- [Kafka, P. \(2026\). Meta’s AI week shows how every company is pushing employees to use AI. Business Insider.](https://www.businessinsider.com/meta-ai-week-employee-training-claude-agents-vibe-coding-2026-3)
- [Blum, S. \(2024\). Klarna Plans to Shut Down SaaS Providers and Replace Them With AI. Inc.](https://www.inc.com/sam-blum/klarna-plans-to-shut-down-saas-providers-and-replace-them-with-ai.html)
- [CX Today. \(2025\). Klarna Didn’t Replace Salesforce — It Replaced Them With Alternative SaaS Apps.](https://www.cxtoday.com/crm/klarna-didnt-replace-salesforce-it-replaced-them-with-alternative-saas-apps/)

```
@article{
  leehanchung,
  author = {Lee, Hanchung},
  title = {The AI Great Leap Forward},
  year = {2026},
  month = {04},
  day = {05},
  howpublished = {\url{https://leehanchung.github.io}},
  url = {https://leehanchung.github.io/blogs/2026/04/05/the-ai-grea
}
```

Brain Food: Credibility is Expensive

SHANE PARRISH (FS) · 04 APR 2026 · [SOURCE](#)

FS | BRAIN FOOD

April 5th, 2026 - #675 - [read online](#) - Free Version

Welcome to Brain Food, your weekly signal in a world full of noise.

Tiny Thoughts

*

One-sided deals are one-time deals.

**

A lot of problems persist because the solution is too simple to take seriously.

Credibility is expensive because the bills never stop.

You pay it in conversations no one overhears, in deals where you leave money on the table, in credit you give away. You pay for it every time you say the hard thing instead of the easy one.

The strange part is that you're paying for years before anyone notices, and you can lose it all in an afternoon.

Insights

*

Jim Rohn on why we set goals:

“ The ultimate reason for setting goals is to entice you to become the person it takes to achieve them. ”

**

Former Prime Minister of the United Kingdom Margaret Thatcher on consensus:

“Consensus ... is the process of avoiding the very issues that have to be solved. ”

Author Ryan Holiday on accomplishment:

“You'll never feel okay by way of external accomplishments. Enough comes from the inside. It comes from stepping off the train. From seeing what you already have, what you've always had. ”

The Knowledge Project

My guest this week is Joe Liemandt, the principal of Alpha School and the founder of Trilogy Software and ESW Capital.

Liemandt dropped out of Stanford to build Trilogy, became the youngest member of the Forbes 400, then vanished from public life for 25 years. But he didn't stop building. Through ESW Capital, he quietly became one of the most prolific acquirers of software businesses in the world.

Now he's back with a \$1 billion bet that AI can help kids learn ten times faster, and that school as we know it is broken.

At Alpha School, students spend two hours a day on AI-driven instruction and score in the top 1% on standardized tests. The rest of the day is devoted to what Liemandt calls life skills: leadership, entrepreneurship, teamwork, and real projects that kids actually care about. There are no lectures, and kids don't move forward until they master the material.

He argues the traditional classroom was designed for a narrow slice of students and wastes everyone else's time. The fix isn't more money or better teachers; it's rebuilding from scratch around mastery, motivation, and AI. The role of a 'teacher' dramatically changes.

This conversation covers everything from sleeping on the floor at Trilogy to being mentored by Jack Welch to the thinking behind Alpha School.

I hope you enjoy this conversation as much as I did.

+ Listen and Learn: [Spotify](#) | [YouTube](#) | [Apple Podcasts](#) | [Web/Transcript](#) | [X](#)

SPONSOR

The Daily Upside*

*Financial headlines can feel overwhelming. Get calm, clear insights—without the fear—from The Daily Upside. Some call it the closest thing to earning your MBA in Cambridge. **Join 1M subscribers who trust The Daily Upside. [Subscribe for free today.](#)***

Thanks for reading. I'll see you next week.

— Shane Parrish

P.S. I can't believe this was [a real stunt](#).

Want to share this edition with a friend? Use this link: <https://fs.blog/brain-food/april-5-2026/>

Upgrade Yourself

Members get access to all my reading highlights, ad-free newsletters and podcasts, hand-edited transcripts, early access, AMAs, and so much more. Join us → [here](#).

You are receiving this email because you subscribed.

Brain Food reaches over 1 million people weekly. Learn more about sponsoring an issue or an episode of The Knowledge Project [here](#).

Overwhelmed by email? No need to unsubscribe. Try a 30 day [break](#). If you want to change your email address, [update your profile](#). Or [unsubscribe](#).

201-854 Bank Street, Ottawa, ON K1S 3W3

The Axios supply chain attack used individually targeted social engineering

SIMON WILLISON'S WEBLOG: ENTRIES · 03 APR 2026 · [SOURCE](#)

The Axios supply chain attack used individually targeted social engineering

The Axios team have published a [full postmortem](#) on the supply chain attack which resulted in a malware dependency going out [in a release the other day](#), and it involved a sophisticated social engineering campaign targeting one of their maintainers directly. Here's Jason Saayman's description of [how that worked](#):

so the attack vector mimics what google has documented here: <https://cloud.google.com/blog/topics/threat-intelligence/unc1069-targets-cryptocurrency-ai-social-engineering>

they tailored this process specifically to me by doing the following:

- they reached out masquerading as the founder of a company they had cloned the companys founders likeness as well as the company itself.
- they then invited me to a real slack workspace. this workspace was branded to the companies ci and named in a plausible manner. the slack was thought out very well, they had channels where they were sharing linked-in posts, the linked in posts i presume just went to the real companys account but it was super convincing etc. they even had what i presume were fake profiles of the team of the company but also number of other oss maintainers.
- they scheduled a meeting with me to connect. the meeting was on ms teams. the meeting had what seemed to be a group of people that were involved.
- the meeting said something on my system was out of date. i installed the missing item as i presumed it was something to do with teams, and this was the RAT.
- everything was extremely well co-ordinated looked legit and was done in a professional manner.

A RAT is a Remote Access Trojan—this was the software which stole the developer’s credentials which could then be used to publish the malicious package.

That’s a *very effective* scam. I join a lot of meetings where I find myself needing to install Webex or Microsoft Teams or similar at the last moment and the time constraint means I always click “yes” to things as quickly as possible to make sure I don’t join late.

Every maintainer of open source software used by enough people to be worth taking in this way needs to be familiar with this attack strategy.

Issue #704

POINTER · 03 APR 2026 · [SOURCE](#)

April 3rd, 2026 | [Read online](#)

Pointer

Friday 3rd March issue is presented by SerpApi



[SerpApi: The Ultimate Sidekick For Your AI Applications](#)

AI is smarter when it can search the web. How are AI companies doing it? Meet the worst-kept secret in the AI industry: the [SerpApi](#) web search API:

👉 Scrape Google and other search engines (including AI overviews, Maps, Shopping, Amazon,

and more)

👉 Integrate with a dead-simple GET request that any agent can make.

👉 Used by NVIDIA, Uber, Adobe, and even the United Nations.

Give it a try now; no sign-up needed.

Explore our Playground

I'm An Introvert. This Is How I Get Myself To Speak Up.

— Wes Kao

tj;dr : I often get readers who ask, “Wes, if I’m not naturally loud, how do I speak up or talk about my accomplishments? How do I gain more visibility, either internally in my company or externally online?”

CareerAdvice

What Kinds Of New Debt Are Teams Accumulating With AI?

— Lizzie Matusov

tj;dr : Dr. Margaret-Anne Storey at the University of Victoria proposes a “triple debt model” for reasoning about software system health, drawing on recent empirical research on developer cognition, AI-assisted development, and software evolution. The paper is a conceptual framework, not a controlled experiment - it synthesizes findings across multiple studies to argue that technical debt alone is an incomplete lens.

Leadership Management

The Ultimate Sidekick for Your AI Applications

tj;dr : AI is smarter when it can search the web. How are AI companies doing it? Meet the worst-kept secret in the AI industry: the SerpApi web search API: (1) Scrape Google and other search engines (including AI overviews, Maps, Shopping, Amazon, and more). (2) Integrate with a dead-simple GET request that any agent can make. (3) Used by NVIDIA, Uber, Adobe, and even the United Nations.

Promoted by SerpApi

Tools

What I Learned From Nearly 1,000 Interviews At Amazon

— Steve Huynh

tl;dr : “A Bar Raiser is a specially trained interviewer whose job is to ensure that every hire raises the average talent level at Amazon. I had veto power over any candidate. I sat on nearly a thousand interview loops across every level from intern to Principal Engineer.” Steven discusses the patterns that emerged.

Leadership Management Hiring

“Whether you think you can or you think you can't, you're right.”

— Henry Ford

No Semicolons Needed

— Terts Diepraam

tl;dr : “I'm making a scripting language called Roto. Like so many programming languages before it, it has the goal of being easy to use and read. Many languages end up making semicolons to delimit or terminate statements optional to that end. I want that too!”

LanguageDesign

IaCConf: How Teams Manage IaC At Scale

tl;dr : Most IaC content focuses on theory. This virtual conference is different. IaCConf brings together engineers sharing what's actually working in production - from scaling to managing complexity across teams. Real lessons from operators and practical patterns you can apply.

Promoted by Spacelift

Event

The Skill Of Using AI Agents Well

tldr : “What tips, tricks, approaches are you using to improve your efficiency when using agent harnesses? I personally focus on Claude Code & Codex CLI mostly for single person software development, but I welcome suggestions for other tools and other areas of use. Ideally you share what you tried and what works for you and I try it myself and see whether it improves my workflow. Here are my discoveries (with level of confidence).”

Agents

Reading Leaked Claude Code Source Code

tldr : “Anthropic accidentally leaked a source map file containing the full, unobfuscated TypeScript source code. 1,897 files around 132,000 LOC. What I found is a piece of software that is surprisingly more opinionated and more paranoid than I expected. Here are some of things that I found worth talking about.”


Claude Design

Architecture Decision Record

— Martin Fowler

tldr : “An Architecture Decision Record (ADR) is a short document that captures and explains a single decision relevant to a product or ecosystem. Documents should be short, just a couple of pages, and contain the decision, the context for making it, and significant ramifications. They should not be modified if the decision is changed, but linked to a superseding decision.”

Null Pointer



Simple Request

Hand-drawn by Manu. [Archives here](#).

Most Popular From Last Issue

[Nobody Is Coming To Save Your Career](#) - Steve Huynh

Notable Links

[Claude Code \(Thread\)](#): Collection of lesser-known features & workflows.

[Hermes](#): Agent that grows with you.

[Openscreen](#): OS alternative to screen studio.

[Open Terminal](#): A computer you can curl.

[Pg textsearch](#): Modern ranked text search for Postgres.

How did you like this issue of Pointer?

1 = Didn't enjoy it all // 5 = Really enjoyed it

[1](#) | [2](#) | [3](#) | [4](#) | [5](#)

