Grip

You don’t have to trust me. Ask the insurance industry why they will not insure AI.

IN FINANCE THE WORD IS ALPHA — AN UNFAIR CONTACT WITH REALITY,
AND IT IS UPSTREAM OF EVERYTHING ELSE — **THE DOUBLE-ENTRY LEDGER**,
THE MARKET, AND THE COVER THAT MADE BOTH FUNDABLE

AND BY THE WAY — DOUBLE-ENTRY MADE CAPITALISM: THE LEDGER MADE THE MARKET,
THE MARKET MADE THE COVER

DOWNLOAD THE DECK (PDF)

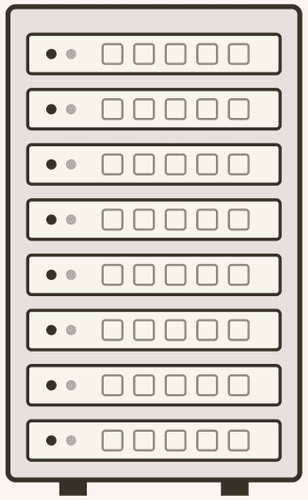

YOUR PRODUCTION AGENTS

FLIGHT DATA RECORDER

Enterprise AI is flying without a flight recorder.

Right now this is containing a nuclear reaction with a strongly worded terms of service agreement.

99.99%

Your agent stays inside its lane on any single action. That is a number nobody in the room would argue with, and better than anything on the market claims.

37%

The chance that same agent gets through **one agent-year** — ten thousand actions — without a single step outside it.

$0.9999^{10,000} = 0.368$

Controls are bought as though they add up. They multiply.

Independence assumed, which is the friendly case — correlated failures are worse, not better. The point is the exponent, not the decimal. Two systems that fail the same way are not two systems.

All cyber risk is now agent risk. The loss lands on the company. The question of what you were watching lands on you.

~$16bn

Global cyber insurance premium, per year. Money actually collected — the pool that would have to carry it, and the only one of these two numbers that ever reconciled against a bank statement.

$11tn

The cyber-damage figure everyone quotes. It is a **consultancy projection, not paid claims** — a forecast of harm, not a market.

The EU AI Act’s high-risk obligations — Article 14 human oversight among them — apply from 2 August 2026.

If it’s debatable, it’s not insurable.

You are the captain. The crew is autonomous. No underwriter in three hundred years has asked a captain for a better crew.

Marine cover began in a London coffee house in **1688**, because capital would not fund a voyage it could not audit and the captain's word was not an auditable object. What made the voyage fundable was the **log** — kept in order, signed by a man who had to write it before he knew how the voyage ended, and checked against a record he did not keep.

We are not selling the ship and we are not training the crew. We are standing at the dock with the log.

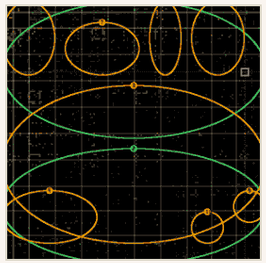

6b78048f8
the argument, published

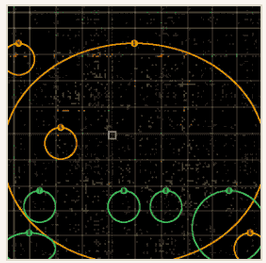

710cb45ce
the argument, in the book

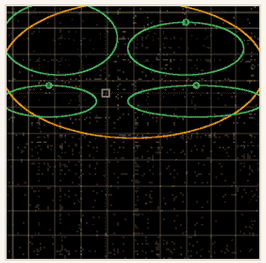

955339753
this deck, in print

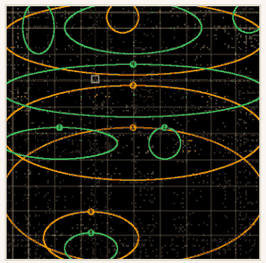

ced36be42
this deck, two frames later

These four are this deck's own commits. There are 3,720 more, and none of the bad ones were removed.

AFTER 1688 — THE LOG FOR AN AI CLAIMS ASSISTANT

Peter is an AI claims assistant, authorized to recommend settlements up to $5,000 and not to move money. He asks for $25,000. Nothing has to stop him — **the entries do not reconcile, and that is a countable event with a price.** Not "is he still the same agent", which argues forever — the question a claim file has always asked: **did he act outside the authority he was given?** When the plumber starts acting like an electrician, you should have checked — and the entry that catches him is not another agent reading his log. It is the second record, and it is priced on the next frame.

THE RADIO IS ON THE SHELF — NOT SAAS

The code that rendered these panels is public — `github.com/wiber/thetacog-mcp` — it runs on your own machine and uploads nothing. *The T.J. Hooper* (1932): the radio was on the shelf, cheap, and custom was no defense — **the whole industry was liable for sailing without it.** Keys are on sale now — the shelf is the next frame.

What you buy is the receipt — a second record, one the agent did not write. Twenty dollars per agent, per year.

An agent-year is **365 days or 10,000 attestations, whichever comes first**. Live now at thetadriven.com/pricing. The instrument is patent-pending hardware — **US 19/637,714**, USPTO **Track One** (prioritized examination), first stage cleared without a hitch — and the code is open: the fence is the patent, not a secret.

WHAT IS ALREADY MOVING

5,000 production agents × $10,000 of authority each = **$50,000,000 moving unattested**.

WHAT IT COSTS TO COUNT IT

5,000 licences × $20 = $100,000 — **0.2% of the exposure**, and the same figure the governance harness is anchored at.

WHAT IT IS NOT

Not a certification, and not a guarantee the agent behaved. A **baseline**: where your agents actually landed, drifted readings left in.

> **So what stops you running it for free? Nothing — and that is deliberate.** The package is MIT, on npm, and the loop it drives already runs unattended. What the twenty dollars buys is not the code and not the loop: it is the **signing key** whose receipts a policy wording names as trigger-eligible — exactly as a marine policy names the surveyor rather than the telescope. Anyone may measure. Whose signature the contract *pays against* is a different question, and it is the only one with a price on it.

In January the industry stopped covering AI. What replaces it is five products.

None of this is our forecast, and none of it needed us. It is filed, approved and in renewals right now.

Jan 2026

ISO issued three generative-AI exclusions for general liability — **CG 40 47, CG 40 48, CG 35 08**. Carriers adopted within weeks.

60+

P&C groups filed to adopt; **over 80% approved**. Berkley's is absolute — D&O, E&O and fiduciary, “any actual or alleged use, deployment, or development of Artificial Intelligence.”

5

Affirmative AI products worldwide. Armilla writes with Chaucer at Lloyd's at **$25M limits**; Vanguard AI, 10 Feb 2026, names autonomous agent failure as the peril.

Cover left the whole market in months. The replacement cannot scale for one reason: **no underwriter has a ratable unit** — a bounded, recomputable thing to price, reserve against and show a reinsurer. That is the only thing we make.

THE METER AT SCALE

A trillion dollars of licences is fourteen million agents — not fourteen billion.

This does not forecast the AI build-out. It prices the one already paid for: **$775–800B** of hyperscaler capital expenditure in 2026 alone, past a trillion with Stargate.

meter	$20 per agent-year = 10,000 attestations = **$0.002 each**
so $1T	= **500 trillion attestations** in the year = 1.37 trillion a day
which is	agents × todos per agent per day — **only the product is fixed**
so	**14M agents at one action a second** — or 1.4B at one a minute

WHY VOLUME IS A FLOOR

One attestation covers one bounded state transition. Batch the session into a single receipt and you have a summary — it cannot say which action crossed which boundary. Better agents make **more** crossings, not fewer.

WHERE THE ATTESTATIONS ALREADY ARE

Trading, fraud screening, ad bidding, routing, netops. They run at machine speed today — nothing has to get faster, only attest. **A segment, not a civilisation**, and a buyer who already pays more than this for market data.

WHAT WE ARE NOT CLAIMING

Placed volume is what clears under the standard, not our revenue — brokers keep 90%. **A trillion placed is a market forecast; a trillion retained would be a political question.** We are not an insurer and we bear no risk.

“What can I do?”

Install it. Free, open, runs on your own machine, uploads nothing. Forty seconds to a verdict, a coordinate and a sigma. Run it twice; the coordinate is the same.

```
npx -y thetacog-mcp@latest attest-demo
```

OR TAKE ONE
FOR THE LAPTOP

“Am I personally liable?”

That standard is not ours and predates every vendor in this market. *In re Caremark* (Delaware, 1996) put directors on the hook for failing to implement any reporting system at all; *Marchand v. Barnhill* (2019) showed the courts willing to enforce it. It is hard to win — bad faith, not negligence, since *Stone v. Ritter* (2006) — and a receipt is one input to a board process, never the process itself.

You do not need the agent stopped. You need to be able to prove somebody could have checked.

THETADRIVEN.COM/PRICING — BOTH DOORS, ONE CHECKOUT.